arXiv:1612.02503v2 [cs.DB] 18 Dec 2016

What do Shannon-type inequalities, submodular width, and disjunctive datalog have to do with one another?

Mahmoud Abo Khamis
LogicBlox, Inc.

Hung Q. Ngo
LogicBlox, Inc.

Dan Suciu
LogicBlox, Inc., &
University of Washington

Abstract

Recent works on bounding the output size of a conjunctive query with functional dependencies and degree bounds have shown a deep connection between fundamental questions in information theory and database theory. We prove analogous output bounds for *disjunctive datalog queries*, and answer several open questions regarding the tightness and looseness of these bounds along the way. The bounds are intimately related to Shannon-type information inequalities. We devise the notion of a "proof sequence" of a specific class of Shannon-type information inequalities called "Shannon flow inequalities". We then show how the proof sequence can be used as symbolic instructions to guide an algorithm called PANDA, which answers disjunctive datalog queries within the size bound predicted. We show that PANDA can be used as a black-box to devise algorithms matching precisely the fractional hypertree width and the submodular width runtimes for queries *with* functional dependencies and degree bounds.

Our results improve upon known results in three ways. First, our bounds and algorithms are for the much more general class of disjunctive datalog queries, of which conjunctive queries are a special case. Second, the runtime of PANDA matches precisely the submodular width bound, while the previous algorithm by Marx has a runtime that is polynomial in this bound. Third, our bounds and algorithms work for queries with input cardinality bounds, functional dependencies, *and* degree bounds.

Overall, our results showed a deep connection between three seemingly unrelated lines of research; and, our results on proof sequences for Shannon flow inequalities might be of independent interest.

1 Introduction

1.1 The problems

In the history of query evaluation in database, logic and constraint satisfaction areas, there are three research threads which have yielded spectacular results recently.

Thread 1: size bound for conjunctive queries. From the seminal work of Grohe and Marx [30], Atserias, Grohe, and Marx [5], and Gottlob, Lee, Valiant and Valiant [27], we now know of a deep connection between the output size bound of a conjunctive query with (or without) functional dependencies (FD) and information theory. In particular, we can derive tight output size bounds by solving a convex optimization problem whose variables are marginal entropies. Briefly, the bound works as follows. Consider a conjunctive query Q represented by a multi-hypergraph $\mathcal{H} = (V, \mathcal{E})$, where $V = [n]$ is identified with the set of variables $A_1, \ldots, A_n$. To each hyperedge $F \in \mathcal{E}$ there is an input relation R_F whose attributes are $(A_i)_{i \in F}$. A function $h : 2^V \to \mathbb{R}_+$ is called *entropic* if there exists a joint distribution on n variables such that $h(F)$ is the marginal entropy of the distribution on the variables in F, for every non-empty set $F \subseteq V$; by convention, $h(\emptyset) = 0$. Let Γ_n^* denote the set of all n-variable entropic functions.[1] Let CC denote the set of "cardinality constraints" of the form $N_F = |R_F|$, obtained from the input database instance. Let FD denote the set of "FD constraints" of the form $X \to Y$, where $\emptyset \subseteq X \subset Y \subseteq V$.[2] From the cardinality- and FD-constraints, we define two classes of set functions:

$$\mathsf{HCC} = \left\{h : 2^V \to \mathbb{R}_+ \mid h(F) \leq \log |N_F|, \forall F \in \mathcal{E}\right\}, \tag{1}$$

$$\mathsf{HFD} = \left\{h : 2^V \to \mathbb{R}_+ \mid h(X) = h(Y), \forall (X, Y) \in \mathsf{FD}\right\}. \tag{2}$$

Let B denote the set of free variables of the conjunctive query Q. ($B = [n]$ if Q is a full conjunctive query.) Then,

$$\log_2 |Q| \leq \sup_{h \in \Gamma_n^* \cap \mathsf{HCC} \cap \mathsf{HFD}} h(B).$$

It is known that Γ_n^* is a cone, but it is not topologically closed [45]. The topological closure $\overline{\Gamma}_n^*$ of Γ_n^* is a convex cone. Hence, we can replace the above bound by

$$\log_2 |Q| \leq \max_{h \in \overline{\Gamma}_n^* \cap \mathsf{HCC} \cap \mathsf{HFD}} h(B).$$

We shall refer to this bound as the *entropic bound*. Recently Gogacz and Toruńczyk [25] showed that the "group-characterizable" entropy function construction from Chan and Yeung [12] is also a construction of a database instance proving that the entropic bound is asymptotically tight.

With the addition of FDs, we do not know how to solve the optimization problem, partly because the convex cone $\overline{\Gamma}_n^*$ is not even polyhedral, due to infinitely many *non-Shannon-type inequalities* [36, 45], To overcome this limitation of the entropic bound, Gottlob et al. [27] proposed that we relax the optimization problem by optimizing over all n-variable polymatroids instead. A polymatroid is a set function $h : 2^V \to \mathbb{R}_+$ which is non-negative, monotone, and submodular, with $h(\emptyset) = 0$ (for definitions, see Sec 2). Let Γ_n denote the set of all n-variable polymatroids, then it is known that $\overline{\Gamma}_n^* \subseteq \Gamma_n$. The relaxed bound, called the *polymatroid bound* is then

$$\log_2 |Q| \leq \max_{h \in \Gamma_n \cap \mathsf{HCC} \cap \mathsf{HFD}} h(B).$$

Every polymatroid h is a $(2^n - 1)$-dimensional real vector. The problem becomes solving a linear program on $O(2^n)$ variables, which is certainly computable.

Without FDs, the polymatroid bound is exactly $\log_2$ of the AGM-bound [5], and it can be computed efficiently in query complexity, because the optimization problem on polymatroids has the same optimal value as a linear program on *modular* functions[3], whose dual is the fractional edge cover of the query's

[1] This is standard notation in information theory [45].

[2] It is possible for $\emptyset \to Y$ to be an FD, if all variables in Y are constants (e.g., when $x = 5, y = 7$ are input predicates).

[3] see the proof of Proposition 6.3

hypergraph. The AGM-bound is (asymptotically) tight [5], in other words the polymatroid bound is tight if there is no FD. In light of this fact, it is natural to wonder the following:

Question 1. Is the polymatroid bound (asymptotically) tight under the presence of functional dependencies?

In a previous work [32], we observed that the aforementioned optimization problem can be extended to handle input relations with "degree constraints". Consider, for example, an input relation $R(A_1, A_2)$, where for every a_1 the following holds: $|\{a_2 \mid (a_1, a_2) \in R\}| \leq d$, where d is some known positive integer. Then, the query is said to satisfy the degree constraint $(\{A_1\}, \{A_1, A_2\}, \log_2 d)$. There have been extensive recent work on exploiting degree constraints in database systems [8–10]. In general, a degree constraint is encoded by the triple $(X, Y, n_{Y|X})$, where $\emptyset \subseteq X \subset Y \subseteq V$, and $n_{Y|X}$ is a positive integer. A monotone set function $h : 2^V \to \mathbb{R}_+$ satisfies the constraint if $h(Y) - h(X) \leq n_{Y|X}$. Note that both cardinality constraints and functional dependency constraints are special cases of degree constraints: a cardinality constraint $h(F) \leq \log N_F$ is a degree constraint of the form $(\emptyset, F, \log_2 N_F)$, and a functional dependency constraint $h(Y) = h(X)$ is a degree constraint of the form $(X, Y, 0)$[4]. Let DC denote the set of all degree constraints, and HDC denote the set of all set functions satisfying the degree constraints:

$$\mathsf{HDC} \stackrel{\text{def}}{=} \left\{h \mid h : 2^V \to \mathbb{R}_+, h(Y) - h(X) \leq d, \forall (X, Y, n_{Y|X}) \in \mathsf{DC}.\right\} \tag{3}$$

The set HDC now subsumes both HCC and HFD. In particular, below are a corresponding version of the entropic and polymatroid bounds for conjunctive queries with degree constraints:

$$\log_2 |Q| \leq \underbrace{\max_{h \in \overline{\Gamma}_n^* \cap \mathsf{HDC}} h(B)}_{\text{entropic bound}} \leq \underbrace{\max_{h \in \Gamma_n \cap \mathsf{HDC}} h(B)}_{\text{polymatroid bound}}. \tag{4}$$

It is natural to ask the following:

Question 2. Is the entropic bound still (asymptotically) tight under degree constraints?

Thread 2: evaluating disjunctive datalog queries. Disjunctive datalog [19] is a powerful extension of datalog with a wide ranging set of applications in AI, databases, logic, and knowledge representation [4]. For example, it was shown in [19] that computing the optimal tour value in a Traveling Salesman Problem, and computing eigenvectors can be expressed in disjunctive datalog but not in datalog with negations. Let $\mathcal{B}$ be a collection of subsets of variables. Associated with each $B \in \mathcal{B}$, there is an atom T_B. The simplest disjunctive datalog query (without negation) is of the form

$$\bigvee_{B \in \mathcal{B}} T_B \coloneq \bigwedge_{F \in \mathcal{E}} R_F \tag{5}$$

where the body of the query is a conjunction as in a typical conjunctive query with hypergraph $\mathcal{H} = (V, \mathcal{E})$. Given a database instance $\mathbf{D}$, a feasible output to the query is a collection of relations $(T_B^{\mathbf{D}})_{B \in \mathcal{B}}$ such that for every tuple $\mathbf{t}$ satisfying the body, there exists a $B \in \mathcal{B}$ such that $\Pi_B(\mathbf{t}) \in T_B^{\mathbf{D}}$. Note that there is no unique feasible output. In light of the entropic and polymatroid bounds for the conjunctive query case, the first natural question for a disjunctive datalog query is

Question 3. What are the entropic and polymatroid size bounds of a minimum feasible output for a disjunctive datalog query (under degree constraints)? Are they tight?

Since a disjunctive datalog query reduces back to a conjunctive query when there is only one atom in the head's disjunction, ideally we would like the bounds to exactly match the bounds (4) when the head of the disjunctive query has only one atom and generalize them to multiple atoms. Once we have derived such a bound, the next natural question is to design disjunctive datalog query evaluation algorithms whose runtimes match that bound. This line of questions is meant to parallel the worst-case optimal join algorithms [38,39,43] matching the AGM-bound, or the CSMA algorithm [32] matching the polymatroid bound (4).

[4]The degree constraint only requires $h(X) \leq h(Y)$, and monotonicity enforces $h(Y) = h(X)$.

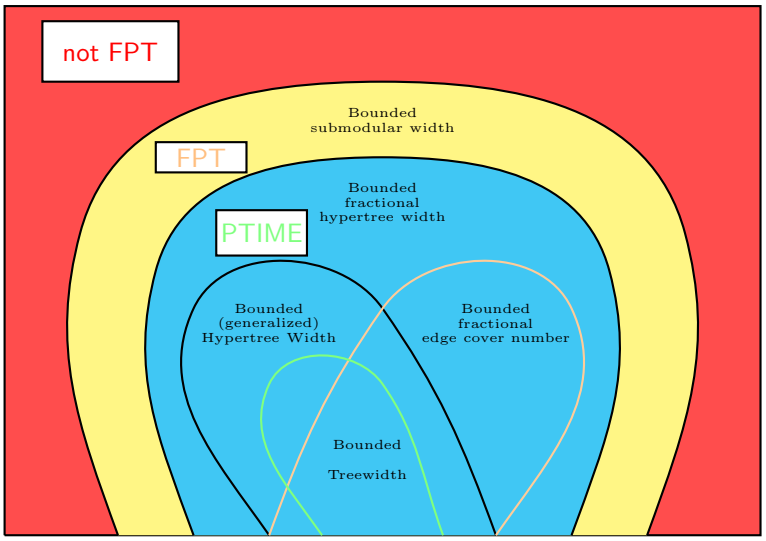

Figure 1: Islands of tractability for conjunctive queries and constraint satisfaction problems.

Question 4. Design a query evaluation algorithm for disjunctive datalog queries whose runtime matches the output size bound(s).

Thread 3: Islands of tractability for conjunctive queries (and CSP problems). There is a long history of research on parameters of a query's hypergraph allowing for efficient evaluation. If the query is fixed, then its size is a constant and the problem is in PTIME; this case is thus not interesting from the tractability standpoint. The aim of this line of research is to study (recursively enumerable) classes of query hypergraphs $\mathcal{C}$ such that *every* query whose hypergraph is in $\mathcal{C}$ is tractable. Two key notions of tractability were studied: polynomial-time (PTIME), and *fixed-parameter tractability* (FPT). Let $\mathcal{I}(\mathcal{C})$ denote the set of conjunctive query instances (or equivalently, CSP instances) whose hypergraph is in the class $\mathcal{C}$. Then, $\mathcal{I}(\mathcal{C})$ is *polynomial-time solvable* if there is a polynomial-time algorithm solving every instance $I \in \mathcal{I}(\mathcal{C})$ in time poly($|I|$) (this is *combined complexity*, where $|I|$ denote the instance size measured in both the data and the query size). The class $\mathcal{I}(\mathcal{C})$ is *fixed-parameter tractable* (FPT, with parameter $\mathcal{H}$, the query's hypergraph) if there is an algorithm solving the instance in time $f(|\mathcal{H}|) \cdot N^c$ for some fixed constant c, where f is any computable function, and N is the data size. In other words, FPT in our context is polynomial-time in the data complexity sense [42].

To define the class $\mathcal{C}$, the following graph parameters are the most common. It was known very early on that if $\mathcal{C}$ is the class of bounded tree-width then $\mathcal{I}(\mathcal{C})$ is in PTIME [23]. Then, after query width [13] came (generalized) hypertree width [28], and fractional hypertree width [30], whose boundedness implies PTIME-classification. In a beautiful paper, Marx showed that when $\mathcal{C}$ is the class of graphs with bounded submodular width [35], then $\mathcal{I}(\mathcal{C})$ is fixed parameter tractable (see also an earlier notion of adaptive width [34]). In fact, Marx also showed that the submodular width is the best one can hope for in terms of being in FPT, unless the exponential time hypothesis fails. Figure 1 explains the important classes of widths known thus far.[5] The hierarchy is strict in the sense that there are (infinite) classes of queries which have bounded submodular width but unbounded fractional hypertree width, bounded fractional hypertree width but unbounded (generalized) hypertree width, and so forth. The reader is referred to Gottlob et al. [26] for a very nice survey on the history of query widths.

The submodular width of a hypergraph $\mathcal{H} = (V, \mathcal{E})$ can be defined briefly using our notation above as follows. Let ED denote the set of *edge-dominated* polymatroids, i.e. the set of polymatroids h such that

[5]Taken from Marx's slides at http://www.cs.bme.hu/~dmarx/papers/marx-weizmann-hypergraph.pdf

$h(F) \leq 1$ for all $F \in \mathcal{E}$.[6] Then, the submodular width of $\mathcal{H}$ (or of the query Q), is defined to be

$$\mathsf{subw}(\mathcal{H}) = \mathsf{subw}(Q) \stackrel{\text{def}}{=} \max_{h \in \Gamma_n \cap \mathsf{ED}} \min_{(T,\chi)} \max_{t \in V(T)} h(\chi(t)),$$

where the max goes over all tree decomposition (T, χ) of $\mathcal{H}$, $T = (V(T), E(T))$ is a tree, and $\chi(t)$ is a bag of the tree decomposition. (See Section 2 for formal definitions.) In the same paper, Marx also derived an algorithm running in data-complexity time $O(\text{poly}(N^{\mathsf{subw}(Q)}))$.[7] Furthermore, he showed that $\mathsf{subw}(\mathcal{H}) \leq \mathsf{fhtw}(\mathcal{H})$, for every hypergraph $\mathcal{H}$, and there are classes of hypergraphs for which $\mathsf{subw}(\mathcal{H})$ is bounded and $\mathsf{fhtw}(\mathcal{H})$ is unbounded. Thus, subw is a strictly better notion of width even if we fix the input query. There are two natural open questions in this line of research:

Question 5. What is the corresponding notion of (fractional hypertree or submodular) widths when the input queries have functional dependencies and, more generally, degree bounds? We shall refer to these widths are the "degree-aware" versions of fractional hypertree and submodular widths, and denote them by da-fhtw and da-subw, respectively.

For the sake of brevity, in what follows and in the remaining of this paper the big-O notation is in data-complexity, hiding a factor that is query-dependent and data-independent. The big-$\tilde{O}$ hides a single log-factor in data-complexity,

From the recent functional aggregate query (or FAQ) framework [31], we know that we know the following: we can compute a query in data-complexity time $\tilde{O}(N^{\mathsf{fhtw}(Q)} + |\mathsf{output}|)$, where Q can be a Boolean conjunctive query, full conjunctive query, count query, or aggregate queries, among others. The notion of "output" depends on the query. For a count or Boolean query, the output size is $O(1)$. For a full conjunctive query, the output size is the number of output tuples, and so on. We would like to have the natural analog of these results in the degree-aware world, leading to the following question:

Question 6. Design an algorithm whose runtime matches the degree-aware notions of width, where Q is either a Boolean, full conjunctive, count, or aggregate query over a single semiring (in the FAQ-sense).

For example, as shown in Example 2.4, 3/2 is the submodular width of the 4-cycle query whose graph is $\mathcal{H} = ([4], \{\{1,2\},\{2,3\},\{3,4\},\{4,1\}\})$. We would like an algorithm solving the Boolean 4-cycle query in time $\tilde{O}(N^{3/2})$, which is non-trivial (see Alon et al. [3]), instead of $\tilde{O}(N^2)$ which is trivial. In Example 6.5, we show that the submodular width of a k-cycle query is $2 - 1/\lceil k/2 \rceil$. An algorithm running within the submodular width time would run in time $\tilde{O}(N^{2-1/\lceil k/2 \rceil})$ – matching the best known algorithm for k-cycle detection (and counting) of Alon, Yuster, and Zwick [3]. Hence, matching the submodular width runtime is not only motivated by islands of tractability questions, but also strongly motivated by the the algorithmic question of finding the fastest algorithm for a *fixed* given query.

1.2 The results

In this paper we answer **all** 6 open questions above in the affirmative. Amazingly, the answers to all 6 questions are very closely related. We devise a coherent framework to study these questions, where in most cases the answer to one question helps pave the way to answering another. We summarize our answers below:

Answer to question 1. The answer is NO, the polymatroid bound is not tight when FDs are present. Worst yet, Theorem 3.1 shows that the gap between the entropic bound and the polymatroid bound can be unbounded for some class of queries. We construct these queries by taking a non-Shannon-type inequality, such as the one from the landmark paper by Zhang and Yeung [46]. This is an inequality satisfied by all functions $h \in \overline{\Gamma}_n^*$, but there is a polymatroid $\bar{h} \in \Gamma_n$ that does not satisfy the inequality. We use the inequality as a guide to construct a database instance giving the counter example; and, by scaling up the instance we can show the infinite gap between the two bounds.

[6] This is just a normalization bound so that the width is data independent.

[7] It is not clear what the exact run time of Marx's algorithm is. His theorem states that it is $O(\text{poly}(N^{\mathsf{subw}(Q)}))$. Our best interpretation of Lemma 4.3 and Lemma 4.5 of [35] is that Marx's algorithm runs in time at least $O(N^{2 \cdot \mathsf{subw}(Q)})$.

Answer to question 2. The answer is YES (Theorem 3.8), the entropic bound is still tight under degree constraints. In fact, we prove a much stronger statement, that the entropic bound for *disjunctive* datalog queries is asymptotically tight, under the presence of FDs and degree constraints. We elaborate more on this answer after the answer to question 3 below.

Answer to question 3. Consider the disjunctive datalog query (5). Given a feasible output $(T_B)_{B\in\mathcal{B}}$, then the quantity $\max_{B\in\mathcal{B}}|T_B|$ is a good measure for the size of this output. In Theorem 3.3, we prove the following

$$\underbrace{\min_{(T_B)_{B\in\mathcal{B}}}}_{\substack{\text{over all}\\ \text{feasible outputs}}} \underbrace{\max_{B\in\mathcal{B}} \log_2 |T_B|}_{\text{log of output size}} \le \underbrace{\max_{h\in\overline{\Gamma}^*_n\cap\mathsf{HDC}} \min_{B\in\mathcal{B}} h(B)}_{\text{entropic bound}} \le \underbrace{\max_{h\in\Gamma_n\cap\mathsf{HDC}} \min_{B\in\mathcal{B}} h(B)}_{\text{polymatroid bound}} \tag{6}$$

where the min on the left-hand-side of the above inequality ranges over all feasible output $(T_B)_{B\in\mathcal{B}}$ of the disjunctive datalog query. It is mathematically pleasing that these bounds *imply* the bounds (4): the generalization is smooth. To get a sense of the non-triviality of the above bound, consider the following disjunctive datalog query (with only cardinality constraints):

$$T_{123}(A_1,A_2,A_3) \vee T_{234}(A_2,A_3,A_4) \text{ :- } R_{12}(A_1,A_2) \wedge R_{23}(A_2,A_3) \wedge R_{34}(A_3,A_4) \wedge R_{14}(A_1,A_4), \tag{7}$$

where all input relations R_{ij} have size at most N. In this case $n=4$, and the constraints HDC say that $h(\{1,2\}), h(\{2,3\}), h(\{3,4\})$, and $h(\{1,4\})$ are all upperbounded by $\log_2 N$. It can be shown that

$$\max_{h\in\Gamma_4\cap\mathsf{HDC}} \min\big\{h(\{1,2,3\}), h(\{2,3,4\}\big\} = \frac{3}{2}\log N, \tag{8}$$

which means there exists a feasible solution T_{123}, T_{234} such that both relations have at most $N^{3/2}$ tuples in them. (In contrast, if the head contained only one atom, say T_{123}, then for some input instances the minimum size of T_{123} would be N^2, and the same lower bound would hold if the head contained a *conjunction* of T_{123} and T_{234}.)

As an answer to Question 2, we also show that the entropic bound is asymptotically tight for disjunctive datalog queries. We use the same approach as that of Gogacz and Toruńczyk [25]: taking the construction of group characterizable entropy functions [12] to build a database instance. Theorem 3.8 is stronger than the main result in [25] in several ways. First, the tightness is on the much more general disjunctive datalog queries. Second, it holds for degree constraints, which generalize both cardinality constraints and FDs. Third, our analysis is more concrete than Gogacz and Toruńczyk because with degree constraints we can not hide behind the bound in the limit (let N go to ∞), because the degree bounds and the input size bounds can be very much different.

Answer to question 4. Let OPT be the optimal polymatroid bound in (6). We present an algorithm called PANDA (for *P*roof-*A*ssisted e*N*tropic-degree *A*ware) algorithm, which computes an answer to a disjunctive datalog query (without negation) in time $O(N\log N + 2^{\mathsf{OPT}}\cdot(\log N)^e)$, where e is a query-dependent number. For example, in the specific case of the disjunctive datalog query (7) above, PANDA runs in time $O(N^{3/2}\log N)$. Note the crucial fact that, as shown in Theorem 3.3 we can always compute a feasible solution to the query so that the sizes of T_{123} and T_{234} are both bounded by $N^{3/2}$, by first computing the set of tuples satisfying the body. However, this strategy takes $\Omega(N^2)$ time in the worst-case.

The seed of PANDA is an idea developed in our recent work [32], where we used a mathematical *proof* of a Shannon-type inequality as symbolic instructions to guide an algorithm. We struggled with three different proofs: chain proof (leading to the chain algorithm, a generalization of the Generic Join algorithm [39] for queries with FDs), submodular proof (leading to the submodularity algorithm), which is based on the Ballister-Bollobás [7] proof of Shearer's inequality, and conditional submodularity proof (leading to the CSMA algorithm) which was our own. None of these proofs were algorithmically satisfactory, for different reasons.

This paper develops the "right" notion of a proof sequence for a general class of Shannon-type inequalities, which we call *Shannon flow inequalities*. The proof sequence instructs PANDA on what to do next. Shannon

flow inequalities are much more general than the notion of "output inequality" developed in [32], which itself already is a generalization of Shearer lemma [14]. Partly for this reason, PANDA is able to solve a much more general problem (of evaluating disjunctive datalog query) than that solved by CSMA (evaluating a full conjunctive query).

To briefly illustrate where a Shannon information inequality comes from and how it is used in PANDA, let us look back at the disjunctive datalog query (7). For any function $h \in \Gamma_4$, we have

$$\begin{aligned}
\min\{h(\{1,2,3\}), h(\{2,3,4\})\} &\leq \frac{1}{2}h(\{1,2,3\}) + \frac{1}{2}h(\{2,3,4\}) \\
&\leq \frac{1}{2}(h(\{1,2\}) + h(\{2,3\}) - h(\{2\})) + \frac{1}{2}(h(\{2,3\}) + h(\{3,4\}) - h(\{3\})) \\
&= \frac{1}{2}(h(\{1,2\}) + h(\{2,3\}) + h(\{3,4\})) + \frac{1}{2}(h(\{2,3\}) - h(\{2\}) - h(\{3\})) \\
&\leq \frac{1}{2}(h(\{1,2\}) + h(\{2,3\}) + h(\{3,4\}).
\end{aligned}$$

(Where the inequalities hold due to submodularity of h.) The input cardinality constraints say that $h(\{1,2\}), h(\{2,3\})$, and $h(\{3,4\})$ are all upperbounded by $\log_2 N$. Hence,

$$\max_{h \in \Gamma_4 \cap \mathsf{HDC}} \min\{h(\{1,2,3\}), h(\{2,3,4\})\} \leq \frac{3}{2}\log_2 N,$$

proving one side of inequality (8). We just gave a proof of the following "Shannon flow inequality":

$$\frac{1}{2}h(\{1,2,3\}) + \frac{1}{2}h(\{2,3,4\}) \leq \frac{1}{2}(h(\{1,2\}) + h(\{2,3\}) + h(\{3,4\})), \quad \forall h \in \Gamma_4. \tag{9}$$

However, in order to turn the inequality into an algorithm, we will need a specific type of proof for it, where we convert the terms on the RHS to the terms of the LHS. For any sets $X \subset Y$, we write $h(Y|X)$ as a short hand for $h(Y) - h(X)$. (If h is entropic, then $h(Y|X)$ is exactly the conditional entropy.) Under this notation, the following two inequalities are just submodularity written in a different way:

$$\begin{aligned}
h(\{1,2\}) + h(\{2,3\}|\{2\}) &\geq h(\{1,2,3\}) \\
h(\{2\}) + h(\{3,4\}) &\geq h(\{2,3,4\}).
\end{aligned}$$

From these two inequality, we can construct a "proof sequence" for inequality (9) as follows:

$$\begin{aligned}
& h(\{1,2\}) + h(\{2,3\}) + (\{3,4\})) \\
\text{(decomposition step)} \quad &= h(\{1,2\}) + h(\{2,3\}|\{2\}) + h(\{2\}) + h(\{3,4\}) \\
\text{(submodularity step)} \quad &\geq h(\{1,2,3\}) + h(\{2\}) + h(\{3,4\}) \\
\text{(submodularity step)} \quad &\geq h(\{1,2,3\}) + h(\{2,3,4\}).
\end{aligned}$$

We think of the terms on the RHS (of (9)) as corresponding to input relations; and, terms of the LHS correspond to *output* relations. Then, PANDA interprets the decomposition step and submodularity steps algorithmically. One way to interpret the decomposition step $h(\{2,3\}) \to h(\{2,3\}|\{2\}) + h(\{2\})$ is to decompose $R_{23}(A_2, A_3)$ into heavy and light values as follows. A value a_2 in the (effective) domain of A_2 is called "heavy" if $|\{a_3 \mid (a_2, a_3) \in R_{23}\}| > \sqrt{N}$, and "light" otherwise. The decomposition step partitions R_{23} into two parts: R^{ℓ}_{23} consists of all tuples (a_2, a_3) for which a_2 is light, and R^{h}_{23} consists of the remaining tuples. Note that the number of heavy A_2-values is bounded by $\sqrt{N}$. After decomposition, the algorithm branches out into two branches, the heavy branch and the light branch. The next step is the submodularity step $h(\{1,2\}) + h(\{2,3\}|\{2\}) \to h(\{1,2,3\})$. This step is applied on the light branch: for every tuple $(a_1, a_2) \in R_{12}$, we search for tuples $(a_2, a_3) \in R^{\ell}_{23}$ and add (a_1, a_2, a_3) to T_{123}. There are at most $N^{3/2}$ tuples added this way. Finally, we apply the last step of the proof sequence, which is the submodularity step $h(\{2\}) + h(\{3,4\}) \to h(\{2,3,4\})$. For every heavy value a_2 (there are, again, $\leq \sqrt{N}$ of those), and for

every $(a_3, a_4) \in R_{34}$ we add (a_2, a_3, a_4) to T_{234}. The size of T_{234} is certainly bounded by $N^{3/2}$, for an overall runtime of $O(N^{3/2})$.

The general PANDA algorithm follows roughly this principle, but is *much* more involved because of functional dependencies and arbitrary degree constraints. The branching factor (due to decomposition steps) is also larger than 2, it is $O(\log N)$, and the number of times we have to branch out becomes the exponent in the $\log N$ factor in the algorithm.

Answer to question 5. To answer this question, we summarized existing width notions using the so-called *maximin* and *minimax* widths in Proposition 6.3. The maximin width is due to Marx [34, 35]. The minimax formulation is probably known to researchers who have worked in this area; however, we couldn't locate a public record where this formulation was written down formally. There is one new observation in Proposition 6.3, which is that the generalized hypertree width turns out to be exactly the (maximin and minimax) subadditive width. Once the maximin and minimax width formulation is in place, adding degree constraints is trivial, leading to the "degree-aware" versions of the fractional hypertree width (da-fhtw) and submodular width (da-subw):

$$\begin{aligned}
\text{da-fhtw}(Q) &\stackrel{\text{def}}{=} \min_{(T,\chi)} \max_{t \in V(T)} \max_{h \in \Gamma_n \cap \mathsf{HDC}} h(\chi(t)) \\
\text{da-subw}(Q) &\stackrel{\text{def}}{=} \max_{h \in \Gamma_n \cap \mathsf{HDC}} \min_{(T,\chi)} \max_{t \in V(T)} h(\chi(t)).
\end{aligned}$$

Note again that degree bounds generalize both cardinality bounds and functional dependencies. We do not normalize the upperbounds to be all 1 as in the submodular width case, because cardinality bounds, FDs, and degree bounds are on different scales. Normalization does not make sense in this general context.

Answer to question 6. We describe our approach assuming the input query is Boolean or full conjunctive. The more general answer on FAQ queries is explained in Section 6.4, which is straightforward.

By precomputing the best tree decomposition minimizing $\max_{t \in V(T)} \max_{h \in \Gamma_n \cap \mathsf{HDC}} h(\chi(t))$, we achieve the da-fhtw bound straightforwardly as follows. For each bag B of the tree decomposition, we run PANDA on the conjunctive query with free variables B, to compute a "bag-relation" T_B. Finally, after semijoin reducing the relations T_B with each of the input relations, we run Yannakakis algorithm to join the tables T_B together, for an overall runtime of $\tilde{O}(2^{\text{da-fhtw}(Q)} \cdot (\log N)^e + |\mathsf{output}|)$. This is summarized in Corollary 6.11.

To achieve the da-subw, a new idea is required. We construct a collection of $O(1)$ disjunctive datalog queries, and run PANDA on each of those within the runtime budget of $O(2^{\text{da-subw}(Q)} \cdot (\log N)^e)$, where e is a number independent of the data. Then, we stitch together the results, using a novel "diagonalization" argument. The overall runtime is $\tilde{O}(2^{\text{da-subw}(Q)} \cdot (\log N)^e + |\mathsf{output}|)$. Note that when there is no FD and degree bounds, and when all input relations have the same size N, then $2^{\text{da-subw}(Q)} = N^{\text{subw}(Q)}$ (because our definition of da-subw does not normalize the bounds as was done in Marx's work). In particular, our result is a smooth generalization of Marx's result. Furthermore, PANDA matches the exponent of the submodular width exactly. The result is proved in Corollary 6.13. The exponent e is dependent on the query size. It is proportional to the length of (the best) proof sequences to Shannon flow inequality.

To explain the diagonalization idea, consider the 4-cycle query

$$C_4(A_1, A_2, A_3, A_4) \text{ :- } R_{12}(A_1, A_2), R_{23}(A_2, A_3), R_{34}(A_3, A_4), R_{41}(A_4, A_1). \tag{10}$$

There are only two non-trivial tree decompositions for it, shown in Figure 2. We run PANDA to answer the following four disjunctive datalog queries:

$$\begin{aligned}
T^1_{123} \vee T^1_{124} &\quad :- \quad R_{12} \wedge R_{23} \wedge R_{34} \wedge R_{14}, \\
T^2_{123} \vee T^1_{234} &\quad :- \quad R_{12} \wedge R_{23} \wedge R_{34} \wedge R_{14}, \\
T^1_{134} \vee T^2_{124} &\quad :- \quad R_{12} \wedge R_{23} \wedge R_{34} \wedge R_{14}, \\
T^2_{134} \vee T^2_{234} &\quad :- \quad R_{12} \wedge R_{23} \wedge R_{34} \wedge R_{14}.
\end{aligned}$$

Figure 2: Two non-trivial non-redundant tree decompositions for 4-cycle

We drop the attributes of the input and output relations to avoid cumbersome notations. It is implicitly understood that relation R_{ij} has attributes $R_{ij}(A_i, A_j)$, and relation T^{ℓ}_{ijk} has attributes $T^{\ell}_{ijk}(A_i, A_j, A_k)$. There are 8 resulting tables. After semijoin reducing each of them with each input relation, we stitch them together to form 8 α-acyclic queries: 4 acyclic queries of the form $T^i_{123} \bowtie T^j_{134}$, $i, j \in [2]$, and 4 acyclic queries of the form $T^i_{124} \bowtie T^j_{234}$, $i, j \in [2]$. We use Yannakakis algorithm to compute the outputs of all these acyclic queries. We claim that C_4 is exactly the union of these outputs. Consider any tuple $\mathbf{a} = (a_1, a_2, a_3, a_4) \in C_4$. Then, from the first disjunctive datalog query, either $(a_1, a_2, a_3) \in T^1_{123}$ or $(a_1, a_2, a_4) \in T^1_{124}$ (or both). Without loss of generality, assume $(a_1, a_2, a_3) \in T^1_{123}$. Similarly, from the last two disjunctive datalog queries, both of the following hold: (i) $(a_1, a_3, a_4) \in T^1_{134}$ or $(a_1, a_2, a_4) \in T^2_{124}$, and (ii) $(a_1, a_3, a_4) \in T^2_{134}$ or $(a_2, a_3, a_4) \in T^2_{234}$. If $(a_1, a_3, a_4) \in T^j_{134}$ for $j \in [2]$, then $\mathbf{a}$ is in the output of the join $T^1_{123} \bowtie T^j_{134}$. If $(a_1, a_3, a_4) \notin T^j_{134}$ for both $j \in [2]$, then $\mathbf{a}$ is in the output of the join $T^2_{124} \bowtie T^2_{234}$. That is the essence of the diagonalization argument. The entire algorithm runs in time $\tilde{O}(N^{3/2} + |\mathsf{output}|$. If we are only computing the Boolean 4-cycle query, then the runtime is $O(N^{3/2} \log N)$, matching that of the AYZ algorithm [3].

2 Background and related work

Throughout the paper, we use the following convention. The non-negative real, rational, and integers are denoted by $\mathbb{R}_+, \mathbb{Q}_+$, and $\mathbb{N}$ respectively. Uppercase A_i denotes a variable/attribute, and lowercase a_i denotes a value in the discrete domain $\mathsf{Dom}(A_i)$ of the variable. For any subset $S \subseteq [n]$, define

$$\begin{aligned} \mathbf{A}_S &= (A_i)_{i \in S}, \\ \mathbf{a}_S &= (a_i)_{i \in S} \in \prod_{i \in S} \mathsf{Dom}(A_i). \end{aligned}$$

In particular, $\mathbf{A}_S$ is a tuple of variables and $\mathbf{a}_S$ is a tuple of specific values with support S. We will sometimes use $\mathbf{t}_S = (t_i)_{i \in S} \in \prod_{i \in S} \mathsf{Dom}(A_i)$ to denote a tuple whose support is S.

We will work on multi-hypergraphs $\mathcal{H} = (V, \mathcal{E})$ (i.e. a hyperedge may occur multiple times in $\mathcal{E}$). A function $f : 2^V \to \mathbb{R}_+$ is called a (non-negative) *set function* on $\mathcal{H}$. A set function f on $\mathcal{H}$ is *modular* if $f(S) = \sum_{v \in S} f(v)$ for all $S \subseteq V$, is *monotone* if $f(X) \le f(Y)$ whenever $X \subseteq Y$, is *subadditive* if $f(X \cup Y) \le f(X) + f(Y)$ for all $X, Y \subseteq V$, and is *submodular* if $f(X \cup Y) + f(X \cap Y) \le f(X) + f(Y)$ for all $X, Y \subseteq V$. Unless specified otherwise, we will only consider *non-negative* and *monotone* set functions f for which $f(\emptyset) = 0$; this assumption will be implicit in the entire paper. When $|V| = n$, let M_n, SA_n, and Γ_n denote the set of all (non-negative and monotone) modular, subadditive, and submodular set functions on V, respectively. (The notation Γ_n is to conform with information theory, to be explained below. Any non-negative modular function is monotone.)

Definition 2.1 (Edge- and vertex-dominated functions)**.** For any finite set U, let $\mathbf{n}_U = (n_u)_{u \in U} \in \mathbb{R}^U_+$ denote a vector of non-negative real numbers, whose coordinates are indexed by U. Let $\mathbf{1}_U$ denote the all-1 vector. Define the following two classes of set functions:

$$\begin{aligned} \mathsf{ED}_{\mathcal{H}}(\mathbf{n}_{\mathcal{E}}) &\stackrel{\text{def}}{=} \{h \mid h : 2^V \to \mathbb{R}_+, h(F) \le n_F, \forall F \in \mathcal{E}\} \\ \mathsf{VD}_{\mathcal{H}}(\mathbf{n}_V) &\stackrel{\text{def}}{=} \{h \mid h : 2^V \to \mathbb{R}_+, h(\{v\}) \le n_{\{v\}}, \forall v \in V\}. \end{aligned}$$

ED stands for *edge-dominated* and VD stands for *vertex-dominated*. Let $\mathsf{ED} \stackrel{\text{def}}{=} \mathsf{ED}_{\mathcal{H}}(\mathbf{1}_{\mathcal{E}})$.

For any two finite sets S and T, we use S^T to denote the collection of all maps $f : T \to S$. Each of these maps f is also viewed as a vector whose coordinates are indexed by members of T and whose coordinate values are members of S.[8]

For the sake of clarify, we focus on a full conjunctive query Q, which is modeled as a (multi-) hypergraph $\mathcal{H} = (V = [n], \mathcal{E})$, where for each hyperedge $F \in \mathcal{E}$ there is an input relation R_F with attributes A_F; the set F is called the *support* of relation R_F. The query is $Q = \bigwedge_{F\in\mathcal{E}} R_F(A_F)$; the hypergraph $\mathcal{H}$ is called the *query hypergraph*. The hypergraph is a multi-hypergraph because for the same hyperedge F there can be multiple relations R_F with the same support F. For each $F \in \mathcal{E}$, let $N_F \stackrel{\text{def}}{=} |R_F|$ (the number of tuples in R_F). Let $N = \max_{F\in\mathcal{E},|R_F|<\infty} N_F$. We will explain later how $|R_F|$ can be ∞ when we deal with queries with FDs, degree bounds, or access restrictions. Section 6.4 explains how we can use our results to deal with Boolean, count, or aggregate queries.

2.1 Queries without FDs nor degree-bounds

2.1.1 Bounds on the worst-case output size

Bounding the worst-case output-size $|Q|$ of a natural join query Q is a well-studied problem. There is a hierarchy of such bounds:

Vertex-size bound A trivial bound is $|Q| \le N^{|V|}$. We shall refer to this bound as the *vertex-size bound* $\mathsf{VB}(Q)$, for lack of better names.

Integral edge cover bound A slightly less trivial bound is the *integral edge cover* bound, described as follows. A vector $(\lambda_S)_{S\in\mathcal{E}} \in \{0,1\}^{\mathcal{E}}$ is called an integral edge cover of $\mathcal{H}$ if it is a feasible solution to the following integer linear program (ILP):

$$\begin{aligned} \min \quad & \sum_{S\in\mathcal{E}} \lambda_S \log_2 N_S \\ \text{s.t.} \quad & \sum_{S:v\in S} \lambda_S \ge 1, \quad \forall v \in V \\ & \lambda_S \in \{0,1\}, \quad \forall S \in \mathcal{E}. \end{aligned} \tag{11}$$

It is easy to show that, for any integral edge cover $(\lambda_F)_{F\in\mathcal{E}}$,

$$|Q| \le \prod_{F\in\mathcal{E}} N_F^{\lambda_F}. \tag{12}$$

(The bound follows from the fact that $Q \subseteq \bigwedge_{\substack{F\in\mathcal{E}\\ \lambda_F=1}} R_F(A_F)$.) Let $(\lambda_F^*)_{F\in\mathcal{E}}$ be an optimal solution to the integral edge cover ILP, then the quantity

$$\rho(Q, (N_F)_{F\in\mathcal{E}}) = \log_2 \left(\prod_{F\in\mathcal{E}} N_F^{\lambda_F^*} \right) \tag{13}$$

is called the log *integral edge cover* bound of $|Q|$. This bound is dependent on the input relations' sizes. Often, to state a bound that is independent of the input size, researchers use a cruder approximation of the bound:

$$|Q| \le \prod_{F\in\mathcal{E}} N_F^{\lambda_F} \le N^{\sum_{F\in\mathcal{E}} \lambda_F}.$$

[8] This is standard combinatorics notation. There are $|S|^{|T|}$ of those maps.

To minimize the upperbound, we solve the following linear program

$$\begin{aligned} \min \quad & \sum_{S \in \mathcal{E}} \lambda_S \\ \text{s.t.} \quad & \sum_{S: v \in S} \lambda_S \geq 1, \quad \forall v \in V \\ & \lambda_S \in \{0, 1\}, \quad \forall S \in \mathcal{E}, \end{aligned}$$

The optimal objective value $\rho(Q)$ of this linear program is called the *integral edge cover number* of the hypergraph $\mathcal{H}$. The integral edge cover number bound is $|Q| \leq N^{\rho(Q)}$. Note that $\rho(Q) \leq |V|$, because to cover all vertices we can pick an arbitrary edge to cover an arbitrary vertex. Hence, the vertex-size bound follows trivially from the integral edge cover number bound.

Fractional edge cover bound Building on ealier works [2, 14, 24, 29], Atserias, Grohe, and Marx [5] observed that the rational relaxation of (11) is still an upper-bound on the query size. In particular, a vector $(\lambda_F)_{F \in \mathcal{E}} \in \mathbb{R}_+^{\mathcal{E}}$ is called a fractional edge cover of $\mathcal{H}$ if it is a feasible solution to the following linear program:

$$\begin{aligned} \min \quad & \sum_{S \in \mathcal{E}} \lambda_S \log_2 N_S \\ \text{s.t.} \quad & \sum_{S: v \in S} \lambda_S \geq 1, \quad \forall v \in V \\ & \lambda_S \geq 0, \quad \forall S \in \mathcal{E}. \end{aligned} \tag{14}$$

From Shearer's lemma [2, 14, 24, 29] they showed that (12) still holds. Let $(\lambda_F^*)_{F \in \mathcal{E}}$ be an optimal solution to the fractional edge cover LP, then the quantity

$$\mathsf{AGM}(Q) = \prod_{F \in \mathcal{E}} N_F^{\lambda_F^*} \tag{15}$$

is called the *AGM-bound* of $|Q|$. Similar to the integral case, we can also use the cruder approximation by solving the LP

$$\begin{aligned} \min \quad & \sum_{S \in \mathcal{E}} \lambda_S \\ \text{s.t.} \quad & \sum_{S: v \in S} \lambda_S \geq 1, \quad \forall v \in V \\ & \lambda_S \geq 0, \quad \forall S \in \mathcal{E}, \end{aligned}$$

The optimal objective value $\rho^*(Q)$ of this linear program is called the *fractional edge cover number* of the hypergraph $\mathcal{H}$. The fractional edge cover number bound is $|Q| \leq N^{\rho^*(Q)}$. Obviously, $\log_2 \mathsf{AGM}(Q) \leq \rho^*(Q) \leq \rho(Q) \leq |V|$.

One remarkable property of the AGM-bound is that it is tight: for any input query Q with prescribed input sizes $N_F, F \in \mathcal{E}$, there exists an input instance satisfying the size bounds such that $|Q| = \tilde{O}(\mathsf{AGM}(Q))$. There are known algorithms [38, 39, 43] with runtime $\tilde{O}(\mathsf{AGM}(Q))$; and thus they are *worst-case optimal*.

2.1.2 Traditional widths of tree decompositions

Worst-case optimal algorithms in the AGM-bound sense are not satisfactory for at least two reasons. First, the output size of an arbitrary natural-join query Q may not be as large as $\mathsf{AGM}(Q)$. Second, many database queries are not natural-join queries; they might be to compute an aggregate from the query, count the output size, or simply determine whether or not the output is empty (Boolean query). Consider, for example, the case when the input natural join query is α-acyclic [20], Yannakakis algorithm [44] can answer the query in

time $\tilde{O}(N + |\mathsf{output}|)$. For instance, suppose $\mathcal{H}$ is a path of length $2k-1$ and all input relations are of size N, then $\mathsf{AGM}(Q) = N^k$. A worst-case optimal algorithm runs in $\tilde{O}(N^k)$-time even if the output is empty. Yannakakis algorithm is much more sensitive to the output size than worst-case optimal join algorithms.

The way to make worst-case optimal join algorithms output-sensitive is to run them on the bags of a tree decomposition of Q instead of running them on Q as a whole. This idea was known in the CSP and graphical model literature since the 80s; the idea is refined with better width parameters by the database research community in the past 15 years or so. We refer the readers to the comprehensive recent survey by Gottlob et al. [26] for historical contexts, technical descriptions, and open problems thereof. In this section, we focus on the main width parameters.

Definition 2.2 (Tree decomposition)**.** A *tree decomposition* of a hypergraph $\mathcal{H} = (V, \mathcal{E})$ is a pair (T, χ), where T is a tree and $\chi : V(T) \to 2^V$ maps each node t of the tree to a subset $\chi(t)$ of vertices such that

- Every hyperedge $F \in \mathcal{E}$ is a subset of some $\chi(t)$, $t \in V(T)$
- For every vertex $v \in V$, the set $\{t \mid v \in \chi(t)\}$ is a non-empty (connected) sub-tree of T.

Somewhat confusingly, the sets $\chi(t)$ are often called the *bags* of the tree decomposition. A tree decomposition is *non-redundant* if no bag is a subset of another. A tree decomposition (T_1, χ_1) is *dominated* by another tree decomposition (T_2, χ_2) if every bag of (T_1, χ_1) is a subset of some bag of (T_2, χ_2). Let $\mathsf{TD}(\mathcal{H})$ denote the set of all non-redundant tree decompositions of $\mathcal{H}$ such that no tree decomposition in $\mathsf{TD}(\mathcal{H})$ is dominated by another tree decomposition in $\mathsf{TD}(\mathcal{H})$. We will also use $\mathsf{TD}(Q)$ to denote $\mathsf{TD}(\mathcal{H})$ where $\mathcal{H}$ is Q's hypergraph. In many cases, we drop the qualifier Q or $\mathcal{H}$ for brevity.

The following property is probably well-known. For completeness, we include a proof sketch in Appendix A.

Proposition 2.3. *Let $n = |V|$, then $|\mathsf{TD}| \le n!$ and the number of bags of a non-redundant tree decomposition is at most n. In particular, $|\mathsf{TD}| = O(1)$ in data complexity.*

A tree decomposition is a generic combinatorial structure used to encode sub-problems of a dynamic programming algorithm for solving the input query. Each bag of a tree decomposition corresponds precisely to a sub-problem. The number of bags in a non-redundant tree decomposition is at most $|V|$ [41]. Hence, the number of sub-problems, given the tree-decomposition is bounded by $|V|$, and the runtime of a dynamic programming algorithm is thus bounded by $|V|$ times the runtime of the worst bag. It is known [1,6,16,31] that a vast number of problems in graphical model inference, database join evaluation, aggregate query computation, constraint satisfaction, and logic problems can be solved using this strategy.

In the context of a natural join query for example, a sub-problem identified by a bag $B = \chi(t)$ of a tree-decomposition (T, χ) can be solved by computing an intermediate relation $I(\mathbf{A}_B)$ whose support is B such that $\Pi_B(Q) \subseteq I(\mathbf{A}_B)$. Once these sub-problems are solved, we run Yannakakis algorithm on the intermediate relations (which form an α-acylic query now) by running semijoin reductions up and down the tree decomposition [22]. After that, the relations on each bag are precisely $\Pi_B(Q)$ and we can then report the output in time $\tilde{O}(|\mathsf{output}|)$. To compute $I(\mathbf{A}_B)$, there are several obvious strategies, each of which leads naturally to a notion of "width" of a tree decomposition.

Tree-width The trivial way to compute $I(\mathbf{A}_B)$ is, for each $i \in \chi(t)$ we take an input relation R_F where $i \in F$, and compute the projection $\Pi_{A_i}(R_F)$. Then, the sub-problem corresponding to B is solved by taking the cross-product of all such projections. Obviously, the time it takes on this sub-problem is $\tilde{O}(N^{|B|})$, and the worst-case runtime of the dynamic programming algorithm is $\tilde{O}(N^{\max_{t \in V(T)} |B|})$. Thus, to apply this strategy one seeks a tree decomposition to minimize the quantity $\max_{t \in V(T)} |\chi(t)|$, i.e. we want to solve the minimax problem $\min_{(T,\chi)} \max_{t \in V(T)} |\chi(t)|$. This is precisely captured by *tree-width*, as was observed in Freuder [23]:

$$\mathsf{tw}(Q) \stackrel{\text{def}}{=} \min_{(T,\chi)} \max_{t \in V(T)} |\chi(t)| - 1. \tag{16}$$

The quantity $\mathsf{tw}(Q)$ is called the *tree-width* of Q (or of the hypergraph $\mathcal{H}$ of Q). The term -1 was added so that the tree-width of a (non-hyper) tree is exactly 1.

Note that the tree-width notion is to minimize the worst-case *vertex-size bound* on the bags. This is not a coincidence: there is a tight connection between worst-case output size bounds and different width parameters of tree decompositions.

(Integral/generalized) hypertree width As observed by Gottlob et al. [26], the next natural strategy to compute $I(\mathbf{A}_B)$ is the following. Consider an integral edge cover $(\lambda_F)_{F\in\mathcal{E}}$ of the bag $B = \chi(T)$; namely, $(\lambda_F)_{F\in\mathcal{E}}$ is a solution to the following integer program:

$$\begin{aligned} \min \quad & \sum_{S\in\mathcal{E}} \lambda_S \log N_S \\ \text{s.t.} \quad & \sum_{S: v\in S} \lambda_S \geq 1, \quad \forall v \in B \\ & \lambda_S \in \{0,1\}, \quad \forall S \in \mathcal{E}. \end{aligned} \tag{17}$$

We compute $I(\mathbf{A}_B)$ by computing the projection $\Pi_B(R_F)$ for all F such that $\lambda_F = 1$, and then compute the join of these projections in time $\prod_{F\in\mathcal{E}} N_F^{\lambda_F}$. Then, the overall runtime is dominated by the worst-case integral edge cover bound on the bags of the tree decomposition. This leads naturally to the *generalized hypertree width* of Q, which is the quantity

$$\mathsf{ghtw}(Q) \stackrel{\text{def}}{=} \min_{(T,\chi)} \max_{t\in V(T)} \rho(\chi(t)). \tag{18}$$

Here, $\rho(\chi(t)) = \rho(B) = \sum_{F\in\mathcal{E}} \lambda_F^*$, where $(\lambda_F^*)_{F\in\mathcal{E}}$ is an optimal solution to the ILP (17) with N_F replaced by N.

Fractional hypertree width As observed in Grohe and Marx [29, 30], one can replace integral by fractional in the Gottlob et al. strategy. The only change is that one has to compute $I(\mathbf{A}_B)$ using a worst-case optimal join algorithm [38, 39, 43]. This observation leads to the notion of *fractional hypertree width*, defined by

$$\mathsf{fhtw}(Q) \stackrel{\text{def}}{=} \min_{(T,\chi)} \max_{t\in V(T)} \rho^*(\chi(t)). \tag{19}$$

Very recently, Fischl et al. [21] showed that, checking whether a given hypergraph has fractional hypertree width or a generalized hypertree width at most 2 is NP-hard, settling two important open questions.

We remark that all of the above definitions of widths capture cruder approximations of what the corresponding algorithms can achieve. The "crudeness" comes from the fact that we replaced N_F by N for all input relations F. The actual runtime of a dynamic-programming algorithm based on worst-case optimal join algorithms is in fact $\tilde{O}\left(N + \min_{(T,\chi)} \max_{t\in V(T)} \mathsf{AGM}(\chi(t))\right)$, which can be much smaller than $\tilde{O}(N^{\mathsf{fhtw}(Q)})$.[9]

2.1.3 Adaptive and submodular widths

The above three width parameters are based on the same idea: we decompose the query into sub-queries in a dynamic-programming algorithm, and the runtime is dominated by the worst bag size bound of the tree decomposition. This strategy has a drawback that once we stick with a tree decomposition we are forced to suffer the worst-case instance for that tree-decomposition. Marx [34, 35] had a wonderful observation: we can partition the data first, and then use a different tree-decomposition for each part of the data, then we can in some cases significantly improve the runtime. This idea leads to the notions of *adaptive width* and *submodular width* of a query. To illustrate Marx's key observation, before presenting the definitions of these widths, let us first consider a simple example.

[9] Consider, for example, the triangle query $Q = R(A_1, A_2) \wedge S(A_2, A_3) \wedge T(A_1, A_3)$ where $|R| = |S| = 1$ and $|T| = N$. The runtime is $\tilde{O}(N)$ instead of $\tilde{O}(N^{3/2})$.

Example 2.4 (Intuition behind submodular width)**.** Consider the following 4-cycle query. (Relation symbols are renamed from (10) for convenience.)

$$C_4(A_1, A_2, A_3, A_4) \text{ :- } R(A_1, A_2), S(A_2, A_3), T(A_3, A_4), W(A_4, A_1). \tag{20}$$

It is easy to see that $\mathsf{ghtw}(C_4) = \mathsf{fhtw}(C_4) = 2$, and there are only two non-trivial tree decompositions for it, shown in Figure 2. The fhtw of the query is 2; because, if we stick with any of the two tree decomposition, we can easily find an input instance for which the worst bag size bound is N^2. For example, suppose we take tree decomposition 1 in the figure, then we can simply take the the database instance in which

$$\begin{aligned} R(A_1, A_2) &= [N] \times \{1\} \\ S(A_2, A_3) &= \{1\} \times [N]. \end{aligned}$$

However, for this *given* input instance, the other tree decomposition has a much better worst bag size bound: the effective domain of A_2 is of size 1, and hence the bags $\{A_2, A_3, A_4\}$ and $\{A_1, A_2, A_4\}$ are both bounded by N. From this observation, Alon, Yuster, and Zwick [3] devised an algorithm to answer the Boolean 4-cycle query above in time $\tilde{O}(N^{3/2})$. The AYZ algorithm works as follows. For any binary relation $R(X, Y)$ of size N, we call a value x "heavy" if $|\{y : (x, y) \in R\}| > \sqrt{N}$, otherwise x is called "light". Clearly there are at most $\sqrt{N}$ heavy x-values. Consider the above 4-cycle query (20). Define

$$\begin{aligned} R_{\mathsf{heavy}}(A_1, A_2) &\stackrel{\text{def}}{=} \{(a_1, a_2) \in R \mid a_1 \text{ is heavy}\} \\ R_{\mathsf{light}}(A_1, A_2) &\stackrel{\text{def}}{=} \{(a_1, a_2) \in R \mid a_1 \text{ is light}\} \\ T_{\mathsf{heavy}}(A_3, A_4) &\stackrel{\text{def}}{=} \{(a_3, a_4) \in T \mid a_3 \text{ is heavy}\} \\ T_{\mathsf{light}}(A_3, A_4) &\stackrel{\text{def}}{=} \{(a_3, a_4) \in T \mid a_3 \text{ is light}\}. \end{aligned}$$

Then, the query C_4 is a *disjoint* union of three sub-queries $C_4 = Q_h \cup Q_{hh} \cup Q_{hl}$, where

$$\begin{aligned} Q_h &\text{ :- } R_{\mathsf{heavy}}(A_1, A_2), S(A_2, A_3), T(A_3, A_4), W(A_4, A_1) \\ Q_{lh} &\text{ :- } R_{\mathsf{light}}(A_1, A_2), S(A_2, A_3), T_{\mathsf{heavy}}(A_3, A_4), W(A_4, A_1) \\ Q_{ll} &\text{ :- } R_{\mathsf{light}}(A_1, A_2), S(A_2, A_3), T_{\mathsf{light}}(A_3, A_4), W(A_4, A_1). \end{aligned}$$

To compute Q_h, we use tree decomposition 1. The intermediate relations $I(A_1, A_3, A_4)$ and $I(A_1, A_2, A_3)$ are both bounded by $N^{3/2}$ because the number of heavy a_1-values is at most $\sqrt{N}$. To compute Q_{lh}, we also use tree decomposition 1, with similar bounds. To compute Q_{ll}, we use tree decomposition 2. The size of the intermediate relation $I(A_1, A_2, A_4)$ is bounded by:

$$\begin{aligned} |I(A_1, A_2, A_4)| &\le |R_{\mathsf{light}}(A_1, A_2) \wedge W(A_4, A_1)| \\ &= \sum_{(a_4, a_1) \in W} |\{a_2 : (a_1, a_2) \in R_{\mathsf{light}}\}| \\ &\le \sum_{(a_4, a_1) \in W} \sqrt{N} \\ &= \sqrt{N} \sum_{(a_4, a_1) \in W} 1 \\ &\le N^{3/2}. \end{aligned}$$

Similarly, $|I(A_2, A_3, A_4)| \le N^{3/2}$. Overall, we can solve the natural join query in time $\tilde{O}(N^{3/2}+Z)$ where Z is the output size. In fact, using the heavy-light partitioning, the AYZ algorithm can compute any aggregate, boolean, or join query on 4-cycle in time $\tilde{O}(N^{3/2} + Z)$ where Z is the output size of the corresponding problem. The data partitioning of Q into Q_h, Q_{lh}, and Q_{ll} is the intuition behind Marx's adaptive width and submodular width.

Definition 2.5 (Adaptive and submodular widths)**.** Recall the notion of edge domination from Defn 2.1. Then, the *adaptive width* $\mathsf{adw}(\mathcal{H})$ of $\mathcal{H}$ and the *submodular width* $\mathsf{subw}(\mathcal{H})$ of $\mathcal{H}$ are defined by

$$\mathsf{adw}(\mathcal{H}) \stackrel{\text{def}}{=} \max_{f \in \mathsf{ED} \cap \mathsf{M}_n} \min_{(T,\chi)} \max_{t \in V(T)} f(\chi(t)), \tag{21}$$

$$\mathsf{subw}(\mathcal{H}) \stackrel{\text{def}}{=} \max_{f \in \Gamma_n} \min_{(T,\chi)} \max_{t \in V(T)} f(\chi(t)). \tag{22}$$

Marx [35] showed that $\mathsf{adw}(\mathcal{H}) \le \mathsf{subw}(\mathcal{H}) \le \mathsf{fhtw}(\mathcal{H})$ $\mathcal{H}$ and $\mathsf{subw}(\mathcal{H}) = O(\mathsf{adw}^4(\mathcal{H}))$, for any $\mathcal{H}$. There are classes of queries for which $\mathsf{subw}(\mathcal{H})$ is bounded by a constant while $\mathsf{fhtw}(\mathcal{H})$ grows with the query size. He also designed an algorithm showing that Boolean conjunctive queries can be solved in data-complexity time $O(\text{poly}(N^{\mathsf{subw}(\mathcal{H})})$.

The definitions of the traditional width parameters and the adaptive and submodular width parameters are somewhat disconnected from one another. In Section 6, we will explain how all these width parameters can be unified using the same information-theoretic framework which can deal with much more general constraints on the input queries. We will also present a simple example showing the unbounded gap between $\mathsf{subw}(\mathcal{H})$ and $\mathsf{fhtw}(\mathcal{H})$ for some class of graphs.

2.2 Join queries with functional dependencies and builtin predicates

The above series of bounds and width parameters were based on a single class of statistics on the input relations: their sizes. In practice we very often encounter queries with functional dependency (FD) information. The FDs come from two main sources: primary keys and builtin predicates (such as $A_1 + A_2 = A_3$, or A_3 is the maximum common prefix of two strings A_1 and A_2). Without loss of generality, we assume the FDs are of the form $X \to Y$, where $X \subset Y \subseteq V$ are subsets of (indices of) variables. (We write $X \to Y$ to denote the FD instead of $\mathbf{A}_X \to \mathbf{A}_Y$ for the sake of brevity.) It is easy to see that AGM-bound is no longer tight in the presence of FDs. For example, the k-clique query has AGM-bound $N^{k/2}$, but if one variable functionally determines the rest of the variables then the output has size at most N.

Gottlob et al. [27] made a key observation allowing us to systematically derive output-size bounds under FDs. To describe their bound and our generalization, we introduce the notion of a "guard" of an FD. Consider a full conjunctive query Q captured by a hypergraph $\mathcal{H} = (V = [n], \mathcal{E})$, where there is an input predicate R_F for every hyperedge $F \in \mathcal{E}$. As before, we assume $N_F = |R_F|$; however, if R_F is a builtin predicate such as $A_1 + A_2 = A_3$, then we set $N_F = +\infty$.

Definition 2.6 (Guard of an FD)**.** An input predicate R_F is said to *guard* an FD $X \to Y$ if the FD is induced by the predicate, where $X \subset Y \subseteq F$.

For example,

- The input predicate $A_3 = 10$ guards the FD $\emptyset \to \{3\}$.
- The equality predicate $A_3 = A_5$ guards two FDs $\{5\} \to \{3,5\}$ and $\{3\} \to \{3,5\}$; so does the input predicate $A_3 + A_5 = 10$.
- The input predicate $A_3 + A_5 = A_6$ guards the FDs $\{3,5\} \to \{3,5,6\}$, $\{5,6\} \to \{3,5,6\}$, and $\{3,6\} \to \{3,5,6\}$.[10]
- If $R(A_1, A_3, A_4, A_5)$ is an input relation where (A_1, A_3) is its primary key, then R guards the FD $\{1,3\} \to \{1,3,4,5\}$.

Let FD denote the set of FDs guarded by the input predicates to Q. The key observation made by Gottlob et al. [27] is the following. Suppose we have already computed the full conjunctive query Q, which has n attributes and $|Q|$ rows. Consider a joint distribution on $\mathbf{A}_V = (A_1, \ldots, A_n)$ constructed by picking $\mathbf{a}_V \in Q$ uniformly with probability $1/|Q|$. For any subset S of V, let $H(S) \stackrel{\text{def}}{=} H[\mathbf{A}_S]$ denote the marginal entropy on the variables $\mathbf{A}_S$. Then, H is a set function $2^V \to \mathbb{R}_+$ satisfying the following conditions:

[10] In the LogicBlox database system, these implied FDs are captured by what we call "primitive rewritings" of input predicates.

- $H(V) = \log_2 |Q|$, because the distribution was uniform.
- $H(F) \leq \log_2 N_F$ for every $F \in \mathcal{E}$, because the support of the marginal distribution on $\mathbf{A}_F$ is a subset of R_F, and the uniform distribution on $\mathbf{A}_F$ has the maximum entropy. (Note that if R_F was a builtin predicate then $N_F = +\infty$ and this constraint trivially holds.)
- $H(Y) = H(X)$ if $(X \to Y) \in \mathsf{FD}$ is an FD guarded by some input predicate.[11]

From these observations, we can derive a hierarchy of output size bounds for queries with FDs in much the same way that there was a hierarchy of output size bounds for queries without FDs described earlier. However, in the FD case it is more natural to describe the bounds from the tightest to the loosest.

Entropic bound A function $h : 2^V \to \mathbb{R}_+$ is called *entropic* if there is a joint distribution on $\mathbf{A}_V$ such that $h(F)$ is precisely the marginal entropy on $\mathbf{A}_F$ [45]. Let Γ_n^* denote the set of all entropic functions on n variables. The above observation leads to the following bound called the *entropic bound*

$$\log_2 |Q| \leq \mathsf{EB}(Q) \stackrel{\text{def}}{=} \max_{h \in \mathsf{HCC} \cap \mathsf{HFD} \cap \overline{\Gamma}_n^*} h(V), \tag{23}$$

where HCC and HFD were defined in (1) and (2). It was shown recently that the entropic bound is tight [25]. The entropic bound suffers from a serious drawback, which is that we do not know how to compute it. The main reason is that Γ_n^* is convex but it is not topologically closed. Its topological closure $\overline{\Gamma}_n^*$ is a convex cone. We would need an infinite number of non-Shannon-type inequalities to characterize even the closure $\overline{\Gamma}_n^*$.

Polymatroid bound A function $h : 2^V \to \mathbb{R}_+$ is called *polymatroid* if it is monotone, submodular, and $h(\emptyset) = 0$. Let Γ_n denote the set of all polymatroids on $V = [n]$. Then, Γ_n is a polyhedral cone [45], in particular it can be described with a finite (exponential in n) number of linear inequalities: those inequalities describe non-negativity, monotonicity, and submodularity. Furthermore, every entropic function is a polymatroid, i.e. $\Gamma_n^* \subseteq \Gamma_n$, but the converse does not hold [45]. This was a breakthrough in information theory [46]. Due to the computational difficulty of dealing with Γ_n^*, Gottlob et al. [27] relax the entropic constraint Γ_n^* to the superset Γ_n:

$$\log_2 |Q| \leq \mathsf{PB}(Q) \stackrel{\text{def}}{=} \max_{h \in \mathsf{HCC} \cap \mathsf{HFD} \cap \Gamma_n} h(V). \tag{24}$$

The polymatroid bound can now be computed by solving the above linear program. (This is also the reason we can use max instead of sup as in the Γ_n^* case.) Recently, Abo Khamis et al. [32] designed an algorithm called CSMA, for queries with FDs which can compute Q in time $O(2^{\mathsf{PB}(Q)} \cdot (\log N)^e)$, where e is an exponent that is independent of the data size. We will have more to discuss on CSMA in what follows.

Chain bound The polymatroid bound's main drawback is that computing it requires solving an exponential-sized linear program. To remedy this problem, Abo Khamis et al. [32] presented the so-called *chain bound* $\mathsf{ChainBound}(Q)$ which (for a fixed chain) can be computed in polynomial-time in query complexity. Furthermore, they designed a corresponding algorithm called the *chain algorithm* running in time $O(2^{\mathsf{ChainBound}(Q)} \log N)$. The chain bound was also shown to be tight on a class of queries whose FD-lattice is a superset of the distributive lattices, and the chain algorithm is worst-case optimal for that class of queries. The chain bound is a generalization of the AGM-bound: when the query has no FD, it is exactly the AGM-bound.

[11] $H(Y) = H(X)$ is equivalent to $H[\mathbf{A}_Y|\mathbf{A}_X] = 0$ where $H[\mathbf{A}_Y|\mathbf{A}_X]$ denotes the conditional entropy.

Subadditive bound Another way to further relax the constraints of the problem is to look only at the class SA_n of non-negative, monotone, and *subadditive set functions* on V. Clearly every polymatroid is subadditive, and thus $\Gamma_n^* \subset \Gamma_n \subset \mathsf{SA}_n$. Naturally, define

$$\mathsf{SAB}(Q) \stackrel{\text{def}}{=} \max_{h \in \mathsf{HCC} \cap \mathsf{HFD} \cap \mathsf{SA}_n} h(V)$$

to be the bound obtained from the polymatroid bound with Γ_n replaced by SA_n. Then, it not hard to show that $\mathsf{SAB}(Q)$ is a generalization of the integral cover bound when there is no FDs:

Proposition 2.7. *When there is no FD, the subadditive bound is precisely the integral edge cover bound:* $\rho(Q, (N_F)_{F \in \mathcal{E}}) = \mathsf{SAB}(Q)$.

(A proof of this simple fact is in the appendix.) Overall, we have the following hierarchy of bounds:

$$\mathsf{EB}(Q) \le \mathsf{PB}(Q) \le \mathsf{ChainBound}(Q) \le \mathsf{SAB}(Q),$$

and every inequality above is strict on some query Q (with FDs). The strictness of the first inequality was not known until our work in this paper.

2.3 Join queries with known degree-constraints

Definition 2.8 (Degree constraints and their guards)**.** A *degree constraint* is a triple $(X, Y, n_{Y|X})$ where $X \subset Y \subseteq V$ and $n_{Y|X} \in \mathbb{R}_+$ is a non-negative real number. A relation R_F is said to guard the degree constraint $(X, Y, n_{Y|X})$ if $X \subset Y \subseteq F$ and for every tuple $\mathbf{t}_X$ we have

$$\deg_{R_F,Y}(\mathbf{t}_X) \stackrel{\text{def}}{=} |\{\mathbf{a}_Y : \mathbf{a}_Y \in \Pi_Y(R_F) \wedge \Pi_X(\mathbf{a}_Y) = \mathbf{t}_X\}| \le 2^{n_{Y|X}}. \tag{25}$$

The quantity on the left hand side is called the *degree of* $\mathbf{t}_X$ with respect to Y in relation R_F.

For example, consider an input relation $R(A_1, A_2, A_3)$ satisfying the following conditions: for every value a_1 in the effective domain of A_1, there are at most d different values $a_2 \in \mathsf{Dom}(A_2)$ such that $(a_1, a_2) \in \Pi_{A_1A_2}(R)$. Then, R guards the degree constraint $(\{A_1\}, \{A_1, A_2\}, \log_2 d)$. Note that a relation may guard multiple degree constraints.

Degree constraints capture additional information we might know about the input predicates. Both the cardinality constraints and the FD constraints are special cases of degree constraints; because the cardinality constraint $|R_F| \le N_F$ *is* the degree constraint $(\emptyset, F, \log_2 N_F)$ guarded by R_F, and the FD constraint $X \to Y$ is the degree constraint $(X, Y, 0)$.

The information-theoretic approach described in the previous section leads naturally to output size bounds for queries with known degree constraints. Let Q be a query and DC denote the set of degree constraints known about Q. (In particular, DC contains all cardinality constraints and FD constraints, plus any degree-bounds one might be able to obtain from the input relations.) Recall the definition of HDC in (3). Then, we straightforwardly obtain the following output size bounds.

(Degree-aware) entropic and polymatroid bounds The output size of Q can be bounded by

$$\log_2 |Q| \le \underbrace{\max_{h \in \mathsf{HDC} \cap \bar{\Gamma}_n^*} h(V)}_{\mathsf{DAEB}(Q)} \le \underbrace{\max_{h \in \mathsf{HDC} \cap \Gamma_n} h(V)}_{\mathsf{DAPB}(Q)} \tag{26}$$

The CSMA algorithm from Abo Khamis et al. [32] can solve a join query Q with known degree constraints in time $O(2^{\mathsf{DAPB}(Q)} \cdot (\log N)^e)$, where e is a data-independent number.

2.4 Aggregate and join queries

So far, we reviewed various bounds and algorithms for conjunctive queries, which basically consist of a join followed by a projection (if the query is not a full conjunctive query). However, there are also aggregate queries, in which the join is followed by some aggregate operation other than a projection [1,6,31,33]. See [31] for more background and related works.

One possible way to define and process such aggregate queries is to formulate them as *functional aggregate queries* [31]. A functional aggregate query (FAQ) has a hypergraph $\mathcal{H} = (V = [n], \mathcal{E})$, whose vertices are the attributes, just like a conjunctive query. However instead of input relations $R_S \subseteq \prod_{i \in S} \mathsf{Dom}(A_i)$ of a conjunctive query, for each hyperedge $S \in \mathcal{E}$, the FAQ has a corresponding *factor* ψ_S, which is a function that maps $\prod_{i \in S} \mathsf{Dom}(A_i)$ to some domain $\mathbf{D}$, i.e. $\psi_S : \prod_{i \in S} \mathsf{Dom}(A_i) \to \mathbf{D}$. While a conjunctive query takes the join of all input relations and projects out some attributes to get a relation over the remaining attributes $F \subseteq [n]$, the FAQ takes the product of all input factors and "aggregates out" some attributes until we are left with a factor φ_F over the remaining attributes $F \subseteq [n]$, i.e. $\varphi_F : \prod_{i \in F} \mathsf{Dom}(A_i) \to \mathbf{D}$. (W.L.O.G we assume $F = [f]$ for some $f \in [n]$.) In particular, the output to FAQ is

$$\varphi_F(\mathbf{a}_F) \stackrel{\text{def}}{=} \bigoplus_{a_{f+1} \in \mathsf{Dom}(A_{f+1})}^{(f+1)} \cdots \bigoplus_{a_n \in \mathsf{Dom}(A_n)}^{(n)} \bigotimes_{S \in \mathcal{E}} \psi_S(\mathbf{a}_S),$$

where for each $i \in [n] \setminus [f]$, the binary operator $\oplus^{(i)}$ either forms a commutative semiring with the product $\otimes$ over the domain $\mathbf{D}$ (i.e. $(\mathbf{D}, \oplus^{(i)}, \otimes)$ is a commutative semiring) or $\oplus^{(i)}$ is the product itself (i.e. $\oplus^{(i)} = \otimes$). The variables in F are called *free* variables. A special and common type of FAQ queries is the case where all the operators $\oplus^{(i)}$ are identical (i.e. $\oplus \stackrel{\text{def}}{=} \oplus^{(f+1)} = \ldots = \oplus^{(n)}$), and $(\mathbf{D}, \oplus, \otimes)$ is a commutative semiring. We refer to this special type of FAQ as FAQ-SS (for FAQ on a single semiring).

In [31], we introduce a width parameter for FAQ queries called the FAQ-width and denoted faqw. The FAQ-width of an FAQ-query generalizes the fractional hypertree width of a conjunctive query. In particular, the FAQ-width of an FAQ-SS-query φ reduces back to the fractional hypertree width of the hypergraph $\mathcal{H}$ of φ when there are no free variables (i.e. $F = \emptyset$) and when all variables are free (i.e. $F = [n]$). We show in [31] how to answer FAQ queries φ with fixed FAQ-width $\mathsf{faqw}(\varphi)$ in time $O(N^{\mathsf{faqw}(\varphi)} + Z)$, where N and Z are the sizes of the input and output using *listing representation* (i.e. the size of a factor ψ_S is taken as the number of tuples $\mathbf{a}_S$ such that $\psi_S(\mathbf{a}_S) \neq \mathbf{0}$). We will see later in Section 6.4 how to extend width bounds/algorithmic results of this paper to FAQ-SS queries.

3 Output size bounds

3.1 The polymatroid size bound is not tight!

Let $I(\mathbf{X}; \mathbf{Y} | \mathbf{Z}) := H(\mathbf{XZ}) + H(\mathbf{YZ}) - H(\mathbf{XYZ}) - H(\mathbf{Z})$ denote the conditional mutual information between random variables $\mathbf{X}$ and $\mathbf{Y}$ conditioned on random variables $\mathbf{Z}$. In a breakthrough paper in information theory, Zhang and Yeung [46] prove that $\overline{\Gamma}_4^* \subsetneq \Gamma_4$ by proving that the inequality

$$2I(A_3; A_4) \leq I(A_1; A_2) + I(A_1; A_3, A_4) + 3I(A_3; A_4 | A_1) + I(A_3; A_4 | A_2)$$

is satisfied by all $h \in \overline{\Gamma}_4^*$, and that there exists $h \in \Gamma_4$ which does not satisfy the inequality. In our language of set functions, the inequality is written as

$$\begin{aligned} & h(\{1\}) + 2h(\{3\}) + 2h(\{4\}) + h(\{1,2\}) + h(\{2,3,4\}) + 4h(\{1,3,4\}) \\ \leq \; & 3h(\{1,3\}) + 3h(\{1,4\}) + 3h(\{3,4\}) + h(\{2,3\}) + h(\{2,4\}). \end{aligned} \tag{27}$$

The more difficult part of their result was to show that inequality (27) is satisfied by all entropic functions h (see also [45]). To prove that there is a polymatroid *not* satisfying it, they only need to exhibit one

polymatroid: for any positive integer n, (it is easy to see that) the following function $\overline{h}$ is a polymatroid not satisfying (27):

$$\overline{h}(F) = \begin{cases} 0 & \text{when } F = \emptyset \\ (|F|+1)\cdot n & \text{for } F \subset [4], |F| \le 2, F \neq \{1,2\} \\ 4\cdot n & \text{for } F \subseteq [4], |F| \ge 3, \text{ or } F = \{1,2\} \end{cases} \tag{28}$$

In particular, $\overline{h} \in \Gamma_4 - \overline{\Gamma}_4^*$. We use this example to show that the polymatroid bound (24) is not tight when there are FDs. In fact, the bound can be arbitrarily worse than the entropic bound. This proves that the polymatroid is not tight under the presence of FDs, resolving an open question in this area of research. In the following theorem, recall the definition of $\mathsf{PB}(Q)$ and $\mathsf{EB}(Q)$ defined in (23) and (24). Note that $\mathsf{PB}(Q)$ and $\mathsf{EB}(Q)$ were defined in the $\log_2$-scale, hence the unbounded difference between the two quantities is an unbounded *ratio* (of $N^{\Omega(s)}$) between the output size bounds.

Theorem 3.1. *There exists a full conjunctive query Q with functional dependencies for which $\mathsf{PB}(Q) > \mathsf{EB}(Q)$. Furthermore, for this query the difference between the two quantities is as large as $\Omega(s \log N)$ where s is the query size.*

Proof. Let n be an arbitrary positive integer. Consider a query called the *Zhang-Yeung* query ZY with the following inputs.

- There are five variables $V = [5]$.
- There is a relation R_F for each $F \subset V = [4]$, $|F| \le 2$, $F \neq \{1,2\}$, where $N_F = |R_F| = 2^{(|F|+1)n}$. There is also a singleton relation $R_{\{5\}}(\mathbf{A}_{\{5\}})$ where $N_{\{5\}} = 2^{2n}$.
- There are the following functional dependencies. First, FDs of the form $\{5,i\} \to V$ for all $i \in [4]$. Second, $\{1,2\} \to V$ is an FD. Third, $F \to V$ for every $F \in \binom{[4]}{3}$.

Due to the FD constraints, every polymatroid $h \in \Gamma_5$ satisfying the FD constraints must satisfy the property that $h(V) = h(\{5,i\}) = h(\{1,2\}) = h(S)$ for every $i \in [4]$ and $S \in \binom{[4]}{3}$. Furthermore, the cardinality constraints state that $h(F) \le (|F|+1)\cdot n$ for $F \subset [4], |F| \le 2, F \neq \{1,2\}$, and $h(\{5\}) \le 2n$.

We now show that $\mathsf{PB}(\mathsf{ZY}) = 4n$. First, for every $h \in \Gamma_5 \cap \mathsf{FD}$, we have $h(V) = h(\{1,2\}) \le h(\{1\}) + h(\{2\}) \le 4n$ due to submodularity. Hence, $\mathsf{PB}(\mathsf{ZY}) \le 4n$. Next, we specify a solution achieving this objective value. All we have to do is to extend $\overline{h}$ in (28) above slightly. Define a function $f \in \Gamma_5$ as follows. (It is trivial to verify that f is indeed a polymatroid.)

$$f(F) = \begin{cases} \overline{h}(F) & F \subseteq [4] \\ \overline{h}([4]) & 5 \in F \text{ and } |F| > 1 \\ 2n & F = \{5\}. \end{cases}$$

We next show that $\mathsf{EB}(\mathsf{ZY}) \le (4 - 1/11)n$. Let h be any function in $\overline{\Gamma}_5^*$ satisfying the FD constraints for this query. From Zhang and Yeung [46] we know h satisfies inequality (27). We have

$$\begin{aligned} 11 \cdot \mathsf{EB}(\mathsf{ZY}) &= 11h([5]) \\ &= h(\{1,5\}) + 2h(\{3,5\}) + 2h(\{4,5\}) + h(\{1,2\}) + h(\{2,3,4\}) + 4h(\{1,3,4\}) \\ &\le h(\{1\}) + 2h(\{3\}) + 2h(\{4\}) + h(\{1,2\}) + h(\{2,3,4\}) + 4h(\{1,3,4\}) + 5h(\{5\}) \\ &\le 3h(\{1,3\}) + 3h(\{1,4\}) + 3h(\{3,4\}) + h(\{2,3\}) + h(\{2,4\}) + 5h(\{5\}) \\ &\le 9n + 9n + 9n + 3n + 3n + 10n \\ &= 43n. \end{aligned}$$

The input size is $N = \Theta(3^n)$. Hence, the difference between $\mathsf{PB}(\mathsf{ZY})$ and $\mathsf{EB}(\mathsf{ZY})$ is $\ge n/11 = \Omega(\log_2 N)$. Finally, to construct a query Q with an amplified gap between the two bounds, consider a query Q which is a cross-product of k independent copies of ZY. In this case, the gap will be $\Omega(k \log_2 N) = \Omega(s \log_2 N)$. □

3.2 Output size bounds for disjunctive datalog queries

Recall that for queries with degree constraints, we have a set DC consisting of elements of the form $(X, Y, n_{Y|X})$ where $X \subset Y \subseteq V$ and $n_{Y|X} \in \mathbb{Q}_+$ is a non-negative rational number. Also, recall the class HDC of set functions defined in (3).

Definition 3.2 (Size bound under degree constraints)**.** Let Q denote a disjunctive datalog query in the form shown in (5) (with degree constraints HDC). Let $\mathcal{F}$ denote a class of non-negative set functions on V. The log-*size-bound of* Q *on* $\mathcal{F}$ is the quantity

$$\mathsf{LogSizeBound}_{\mathcal{F}}(Q) \stackrel{\text{def}}{=} \sup_{h \in \mathcal{F}} \min_{B \in \mathcal{B}} h(B). \tag{29}$$

Theorem 3.3. *For any disjunctive datalog query Q of the form (5), the following holds:*

$$\min_{(T_B)_{B\in\mathcal{B}}} \max_{B\in\mathcal{B}} \log_2 |T_B| \le \mathsf{LogSizeBound}_{\overline{\Gamma}_n^* \cap \mathit{HDC}}(Q) \le \mathsf{LogSizeBound}_{\Gamma_n \cap \mathit{HDC}}(Q), \tag{30}$$

where the minimum on the left-most quantity ranges over all feasible outputs $(T_B)_{B\in\mathcal{B}}$ to the disjunctive datalog query Q.

Proof. Since $\overline{\Gamma}_n^* \subseteq \Gamma_n$, the second inequality is trivial. To show the first inequality, let T denote the set of all tuples $\mathbf{t}$ satisfying the body of Q; construct a set $\overline{T}$ of tuples as follows. We scan though tuples $\mathbf{t} \in T$ one at a time, and either add $\mathbf{t}$ to $\overline{T}$ or ignore $\mathbf{t}$. To decide whether to add $\mathbf{t}$ to $\overline{T}$, we also keep a collection of tables $(\overline{T}_B)_{B\in\mathcal{B}}$. These tables shall form a feasible output to the disjunctive datalog query Q. Initially $\overline{T}$ and all the $\overline{T}_B$ are empty. Consider the next tuple $\mathbf{t}$ taken from T. If $\Pi_B(\mathbf{t}) \in \overline{T}_B$ for *any* $B \in \mathcal{B}$, then we ignore $\mathbf{t}$. Otherwise, we add $\Pi_B(\mathbf{t})$ to $\overline{T}_B$ for *every* $B \in \mathcal{B}$, and add $\mathbf{t}$ to $\overline{T}$. In the end, obviously the collection $(\overline{T}_B)_{B\in\mathcal{B}}$ is a feasible output to the disjunctive query. Furthermore, by construction the tuples $\mathbf{t} \in \overline{T}$ satisfy the property that: for every two different tuples $\mathbf{t}, \mathbf{t}' \in \overline{T}$, every $B \in \mathcal{B}$, we have $\Pi_B(\mathbf{t}) \neq \Pi_B(\mathbf{t}')$ and both are in $\overline{T}_B$.

Now, construct a joint probability distribution on n variables by picking uniformly a tuple from $\overline{T}$. Let $\overline{h}$ denote the entropy function of this distribution, then

$$\log_2 |\overline{T}| = \overline{h}(V) = \overline{h}(B), \forall B \in \mathcal{B}.$$

Furthermore, $\overline{h} \in \overline{\Gamma}_n^*$ by definition, and $\overline{h} \in \mathsf{HDC}$ because T was a set of tuples satisfying all input degree constraints. Consequently,

$$\begin{aligned}
\min_{(T_B)_{B\in\mathcal{B}}} \max_{B\in\mathcal{B}} \log_2 |T_B| &\le \max_{B\in\mathcal{B}} \log_2 |\overline{T}_B| \\
&= \max_{B\in\mathcal{B}} \overline{h}(B) \\
&= \min_{B\in\mathcal{B}} \overline{h}(B) \\
&\le \max_{h \in \overline{\Gamma}_n^* \cap \mathsf{HDC}} \min_{B\in\mathcal{B}} h(B) \\
&= \mathsf{LogSizeBound}_{\overline{\Gamma}_n^* \cap \mathsf{HDC}}(Q).
\end{aligned}$$

□

A proof of the following simple proposition can be found in Appendix D, where we prove it along with Proposition 6.3. Note that the proposition discusses conjunctive query, where something a bit stronger can be said about the size bounds: some bounds collapse part of the hierarchy of function classes: $\mathsf{M}_n \subset \overline{\Gamma}_n^* \subset \Gamma_n \subset \mathsf{SA}_n$. (Review notation in Section 2.)

Proposition 3.4. *Let Q be a conjunctive query with no FDs whose hypergraph is $\mathcal{H} = (V, \mathcal{E})$. Then the followings hold:*

$$\log_2 \mathsf{VB}(Q) = \mathsf{LogSizeBound}_{\mathcal{F} \cap \mathsf{VD}(\log_2 N \cdot \mathbf{1}_V)}(Q), \quad \forall \mathcal{F} \in \{\mathsf{M}_n, \overline{\Gamma}^*_n, \Gamma_n, \mathsf{SA}_n\} \tag{31}$$

$$\rho(Q, (N_F)_{F \in \mathcal{E}}) = \mathsf{LogSizeBound}_{\mathsf{SA}_n \cap \mathsf{ED}((\log_2 N_F)_{F \in \mathcal{E}})}(Q) \tag{32}$$

$$\rho^*(Q) \cdot \log_2 N = \mathsf{LogSizeBound}_{\mathcal{F} \cap \mathsf{ED}(\log_2 N \cdot \mathbf{1}_{\mathcal{E}})}(Q), \quad \forall \mathcal{F} \in \{\mathsf{M}_n, \overline{\Gamma}^*_n, \Gamma_n\} \tag{33}$$

$$\log_2 \mathit{AGM}(Q) = \mathsf{LogSizeBound}_{\mathcal{F} \cap \mathsf{ED}((\log_2 N_F)_{F \in \mathcal{E}})}(Q), \quad \forall \mathcal{F} \in \{\mathsf{M}_n, \overline{\Gamma}^*_n, \Gamma_n\} \tag{34}$$

Regarding the AGM-bound, the fact that the size bound on Γ_n is equal to the size bound on M_n allows us to compute the bound efficiently (polynomial time in query complexity). By contrast, computing the integral edge cover bound is NP-hard in query complexity, despite being the worse bound: over the more relaxed function class SA_n. From the above proposition and the chains of inclusion $M_n \subset \overline{\Gamma}^*_n \subset \Gamma_n \subset \mathsf{SA}_n$ and $\mathsf{ED}(\mathbf{1}_{\mathcal{E}}) \subseteq \mathsf{VD}(\mathbf{1}_V)$, the following follows trivially

Corollary 3.5. *Let Q be a conjunction query (with no FDs nor degree constraints) whose hypergraph is $\mathcal{H}$, then*

$$\mathsf{VB}(Q) \geq 2^{\rho(Q, (N_F)_{F \in \mathcal{E}})} \geq \mathit{AGM}(Q). \tag{35}$$

And, the gap between any two consecutive entries in the above two series is unbounded.

3.3 The entropic size bound is asymptotically tight under degree constraints for disjunctive datalog queries

Recently, Gogacz and Toruńczyk [25] showed that $\mathsf{LogSizeBound}_{\overline{\Gamma}^*_n \cap \mathsf{HCC} \cap \mathsf{HFD}}(Q)$ is an asymptotically tight upper bound for full *conjunctive* queries Q with FDs. This section proves a generalization of their result. We show that $\mathsf{LogSizeBound}_{\overline{\Gamma}^*_n \cap \mathsf{HDC}}(Q)$ is asymptotically tight for *disjunctive datalog* queries Q with given degree constraints HDC. The proof is based on the same observation made by Gogacz-Toruńczyk, which is that one can take the *group characterizable entropy function* from Chan and Yeung [12] and turn it into a database instance.

Three minor technicalities we need to nail down to push the proof through are the following. First, we need to make sure that the database instance so constructed satisfies the given degree constraints (and the FD constraints). Second, it takes a bit more care to define what we meant by "asymptotically tight" in the case of queries with degree bounds. When the input query has only input relation cardinality bounds and functional dependencies, one can set all input relations to be of the same cardinality N and let N go to infinity. The tightness of the bound is in the exponent α (if the bound was N^α). This was the result from Gogacz and Toruńczyk. When the input query has degree bounds, it does not make sense to set all degree bounds (and cardinality bounds) to be N: we want a finer level of control over their relative magnitudes. Third, unlike in the full conjunctive query case where there is only one output; in the disjunctive datalog query case there are multiple output and we have to show that *any* feasible output to the query must have size asymptotically the same as the worst-case entropic bound.

Definition 3.6 (Database instance from group system)**.** Let G be a finite group and $G_1, \ldots, G_n$ be n subgroups of G. A database instance associated with this group system $(G, G_1, \ldots, G_n)$ is constructed as follows. There are n attributes $A_1, \ldots, A_n$. The domain of attribute A_i is the left coset space G/G_i, i.e. the collection of all left cosets $\{gG_i \mid g \in G\}$ of the subgroup G_i. For every $F \subseteq [n]$, define a relation R_F by

$$R_F \stackrel{\text{def}}{=} \{(gG_i)_{i \in F} \mid g \in G\}.$$

In other words, R_F is the set of tuples $\mathbf{a}_F = (a_i)_{i \in F}$ on attributes $\mathbf{A}_F$, where $a_i = gG_i$ for some $g \in G$.

Given a group system $(G, G_1, \ldots, G_n)$, for any $\emptyset \neq F \subseteq [n]$, define $G_F \stackrel{\text{def}}{=} \bigcap_{i \in F} G_i$; and, define $G_\emptyset = G$. Note that G_Z is a subgroup of G_Y for all $Y \subseteq Z \subseteq [n]$. The following observation is a slight generalization of Proposition 4 from Chan [11]. Recall the notion of degree defined in (25), we prove the following lemma.

Lemma 3.7. *Let* $\mathbf{D}$ *be a database instance associated with the group system* $(G, G_1, \ldots, G_n)$*. Let* $\emptyset \subseteq Z \subset Y \subseteq [n]$ *be two sets of attributes (*Z *can be empty). Then, for any tuple* $\mathbf{a}_Z \in R_Z$*:*

$$\deg_{R_Y,Y}(\mathbf{a}_Z) = \frac{|G_Z|}{|G_Y|}.$$

In particular, $Z \to Y$ *is an FD satisfied by the database instance iff* G_Z *is a subgroup of* $\bigcap_{i\in Y-Z} G_i$*.*

Proof. For any $F \subseteq [n]$, two group elements $b, c \in G$ define the same tuple $\mathbf{a}_F$ in R_F, namely $(cG_i)_{i\in Z} = (bG_i)_{i\in Z}$, if and only if $b^{-1}c \in \bigcap_{i\in F} G_i = G_F$, which is equivalent to $b = cg$ for some $g \in G_F$.

Now suppose $\mathbf{a}_Z \in R_Z$ was defined by an element $b \in G$, i.e. $\mathbf{a}_Z = (a_i)_{i\in Z} = (bG_i)_{i\in Z}$. Consider any two group elements $c_1, c_2 \in G$ which define the same $\mathbf{a}_Z$ tuple but different $\mathbf{a}_Y$ tuples. First, $c_1, c_2 \in G$ define the same $\mathbf{a}_Z$ iff there are two elements $g_1, g_2 \in G_Z$ such that $c_1 = bg_1$ and $c_2 = bg_2$. Second, c_1 and c_2 define different $\mathbf{a}_Y$ tuples iff $c_1^{-1}c_2 \notin G_Y$, which is equivalent to $g_1^{-1}b^{-1}bg_2 = g_1^{-1}g_2 \notin G_Y$, which in turn is equivalent to $g_1G_Y \neq g_2G_Y$. Thus, the degree of $\mathbf{a}_Z$ in R_Y is precisely the index of the subgroup G_Y in the group G_Z, which is equal to $|G_Z|/|G_Y|$.

If $Z \to Y$ is a functional dependency, then from the fact that G_Y is a subgroup of G_Z and the inequality $|G_Z|/|G_Y| \leq 1$, we conclude that G_Z has to be a subgroup of $\bigcap_{i\in Y-Z} G_i$. The converse holds trivially. □

For a given set of degree constraints HDC, let $\mathsf{HDC} \times k$ denote the same set of constraints but all the degree bounds are multiplied by k. The constraints $\mathsf{HDC} \times k$ are called the "scaled up" degree constraints.

Theorem 3.8. *Let* Q *be the* disjunctive datalog *query of the form* (5) *with degree constraints* HDC *for which* $\mathsf{LogSizeBound}_{\overline{\Gamma}^*_n \cap \mathsf{HDC}}(Q) \geq 1$*. For any* $\epsilon > 0$*, there exists a scale factor* k *and a database instance* $\mathbf{D}$ *satisfying the scaled-up degree constraints* $\mathsf{HDC} \times k$ *for which the following holds:*

$$\min_{(T_B)_{B\in\mathcal{B}}} \max_{B\in\mathcal{B}} \log_2 |T_B^{\mathbf{D}}| \geq (1-\epsilon) \cdot \mathsf{LogSizeBound}_{\overline{\Gamma}^*_n \cap \mathsf{HDC}\times k}(Q)$$

where the min *on the left-hand-side is over all feasible outputs to the query* Q*.*

Proof. Let $H = \operatorname{argmax}_{h\in\overline{\Gamma}^*_n\cap\mathsf{HDC}} \min_{B\in\mathcal{B}} h(B)$. Note that, $\overline{\Gamma}^*_n \cap \mathsf{HDC}$ is a convex cone, and thus kH is an optimal solution to the problem $\max_{h\in\overline{\Gamma}^*_n\cap\mathsf{HDC}\times k} \min_{B\in\mathcal{B}} h(B)$. Following Chan and Yeung [12] without loss of generality we assume that there is a distribution on n random variables $A_1, \ldots, A_n$ such that all their domains $\mathsf{Dom}(A_i)$ are discrete and finite, and that H is the joint entropy all of whose marginal entropies are rational numbers. (In general, just as in Chan and Yeung [12], if H is not rational, or is not an entropy over discrete random variables of finite domains, we construct a sequence of distributions over discrete and finite domains whose entropies tend to H.)

For any tuple $\mathbf{a}_F \in \prod_{i\in F} \mathsf{Dom}(A_i)$, we use $p(\mathbf{a}_F)$ to denote $\Pr[\mathbf{A}_F = \mathbf{a}_F]$, i.e. p is the probability mass function. Let d denote the minimum common denominator of all probabilities $p(\mathbf{a}_{[n]})$, over all tuples $\mathbf{a}_{[n]}$ for which $p(\mathbf{a}_{[n]}) > 0$. Let lg denote $\log_2$. Set δ such that $\epsilon = \frac{2\delta}{1+\delta}$, and let r be a multiple of d sufficiently large so that $d \lg e + \frac{d}{2}\lg d \leq \delta r$ and $\lg r \geq d \lg(e^2)$. Construct a matrix $\mathbf{M}_r$ whose columns are in $\prod_{i=1}^n \mathsf{Dom}(A_i)$ where column $\mathbf{a}_{[n]}$ appears precisely $r \cdot p(\mathbf{a}_{[n]})$ times.

The following construction of a group system is from Chan and Yeung [12]. Let G^r denote the group of permutations of columns of $\mathbf{M}_r$. This group acts on the rows of $\mathbf{M}_r$. For $i \in [n]$, let G_i^r be the subgroup of G^r that fix the ith row of $\mathbf{M}_r$, i.e G^r is the stabilizer subgroup of G^r with respect to the ith row of $\mathbf{M}_r$. Note that $|G_F^r| = \prod_{\mathbf{a}_F}(r \cdot p(\mathbf{a}_F))!$, where the product is over all vectors $\mathbf{a}_F$ which have positive probability mass. Consider the database instance $\mathbf{D}^r$ associated with the group system $(G^r, G_1^r, \ldots, G_n^r)$. We claim that this database instance satisfies all degree constraints $\mathsf{HDC} \times k$ for scale factor $k = r(1+\delta)$.

First, we verify that $\mathbf{D}^r$ satisfies all the functional dependency constraints (i.e. the degree upperbounds of 0). Let $Z \to Y$ be an FD. Then, $H[\mathbf{A}_Y|\mathbf{A}_Z] = 0$, which means for every tuple $\mathbf{a}_Z$ with positive probability mass, the tuple $\mathbf{a}_{Y-X}$ is completely determined. In particular, there cannot be two vectors $\mathbf{a}_Y \neq \mathbf{a}'_Y$ with positive probability mass for which $\Pi_X(\mathbf{a}_Y) = \Pi_X(\mathbf{a}'_Y)$. Thus, if a permutation fixes all the rows in Z, then it also fixes all the rows in $Y - Z$ of the matrix $\mathbf{M}_r$. This means G_Z is a subgroup of $\bigcap_{i\in Y-Z} G_i$. From Lemma 3.7, the associated database instance satisfies the FD.

Second, we verify the higher-order degree constraints (with degree bounds $n_{Y|Z} = \log_2 N_{Y|Z} \geq 1$). From definition of probability mass, entropy, and the monotonicity of the $\lg = \log_2$ function, we know the obvious facts that

$$\begin{aligned}
\sum_{\mathbf{a}_F} p(a_F) &= 1, \\
-\sum_{\mathbf{a}_F} p(\mathbf{a}_F) \lg p(\mathbf{a}_F) &= H[\mathbf{A}_F], \\
\sum_{\mathbf{a}_F} \lg(c \cdot p(\mathbf{a}_F)) &\leq d \cdot \lg\left(\sum_{\mathbf{a}_F} c \cdot p(\mathbf{a}_F)\right) = d \cdot \lg c. \\
\sum_{\mathbf{a}_F} \lg\left(\frac{1}{p(\mathbf{a}_F)}\right) &\leq d \cdot \lg d.
\end{aligned}$$

In the above and henceforth in this section, $\sum_{\mathbf{a}_F}$ denotes the sum over all positive-mass vectors $\mathbf{a}_F$. From Stirling approximation[12] we have

$$\begin{aligned}
\lg |G_F^r| &= \lg \prod_{\mathbf{a}_F} (r \cdot p(\mathbf{a}_F))! \\
&\leq \sum_{\mathbf{a}_F} \left(\lg(e) + r \cdot p(\mathbf{a}_F) \lg (r \cdot p(\mathbf{a}_F)) + \frac{1}{2} \lg(r \cdot p(\mathbf{a}_F)) - r \cdot p(\mathbf{a}_F) \lg e \right) \\
&= r \lg(r/e) - rH[\mathbf{A}_F] + \sum_{\mathbf{a}_F} \left(\lg e + \frac{1}{2} \lg(r \cdot p(\mathbf{a}_F)) \right)
\end{aligned}$$

Similarly, $\lg |G_F^r| \geq r \lg(r/e) - rH[\mathbf{A}_F] + \sum_{\mathbf{a}_F} \left(\lg \sqrt{2\pi} + \frac{1}{2} \lg(r \cdot p(\mathbf{a}_F)) \right)$. We now use Stirling approximation as above and Lemma 3.7 to bound

$$\begin{aligned}
\log_2 \frac{|G_Z^r|}{|G_Y^r|} &= \log_2 \frac{\prod_{\mathbf{a}_Z} (r \cdot \Pr[\mathbf{A}_Z = \mathbf{a}_Z])!}{\prod_{\mathbf{a}_Y} (r \cdot \Pr[\mathbf{A}_Y = \mathbf{a}_Y])!} \\
&\leq \left[r \lg(r/e) - rH[\mathbf{A}_Z] + \sum_{\mathbf{a}_Z} \left(\lg e + \frac{1}{2} \lg(r \cdot p(\mathbf{a}_Z)) \right) \right] \\
&\quad - \left[r \lg(r/e) - rH[\mathbf{A}_Y] + \sum_{\mathbf{a}_Y} \left(\lg \sqrt{2\pi} + \frac{1}{2} \lg(r \cdot p(\mathbf{a}_Y)) \right) \right] \\
&= rH[\mathbf{A}_Y|\mathbf{A}_Z] + \sum_{\mathbf{a}_Z} \left(\lg e + \frac{1}{2} \lg(r \cdot p(\mathbf{a}_Z)) \right) - \sum_{\mathbf{a}_Y} \left(\lg \sqrt{2\pi} + \frac{1}{2} \lg(r \cdot p(\mathbf{a}_Y)) \right) \\
&\leq rH[\mathbf{A}_Y|\mathbf{A}_Z] + d \lg e + \frac{1}{2} \sum_{\mathbf{a}_Y} \lg \frac{1}{p(\mathbf{a}_Y)} \\
&\leq rH[\mathbf{A}_Y|\mathbf{A}_Z] + d \lg e + \frac{d}{2} \lg d \\
&\leq rH[\mathbf{A}_Y|\mathbf{A}_Z] + \delta r \\
&\leq r n_{Y|Z} + \delta r \\
&\leq r(1+\delta) n_{Y|Z} \\
&= k \cdot n_{Y|Z}.
\end{aligned}$$

Third, let $(T_B^{\mathbf{D}})_{B \in \mathcal{B}}$ be any feasible output to the query Q on the database instance $\mathbf{D}$ satisfying the scaled

[12] $\lg(\sqrt{2\pi}) + n \lg n + \frac{1}{2} \lg n - n \lg e \leq \lg n! \leq \lg(e) + n \lg n + \frac{1}{2} \lg n - n \lg e$

degree constraints $\mathsf{HDC} \times k$. We are to show that

$$\max_{B\in\mathcal{B}} \lg |T_B^{\mathbf{D}}| \geq (1-\epsilon)\cdot \mathsf{LogSizeBound}_{\overline{\Gamma}_n^* \cap \mathsf{HDC}\times k}(Q) = (1-\epsilon)\cdot k \cdot \min_{B\in\mathcal{B}} H[\mathbf{A}_B]. \tag{36}$$

To bound $\max_{B\in\mathcal{B}} |T_B^{\mathbf{D}}|$, we reason as follows. For any $B \in \mathcal{B}$, we say that $\mathbf{t}_B \in R_B$ "covers" a tuple $\mathbf{a} \in Q$ whenever $\Pi_B(\mathbf{a}) = \mathbf{t}_B$. From Lemma 3.7, for every $B \in \mathcal{B}$, and every tuple $\mathbf{t}_B \in R_B$, there are precisely $\frac{|G_B^r|}{|G_{[n]}^r|}$ tuples $\mathbf{a} \in Q$ covered by $\mathbf{t}_B$. The size of R_B is exactly $\frac{|G_\emptyset^r|}{|G_B^r|}$ and the size of Q (i.e. $R_{[n]}$) is exactly $\frac{|G_\emptyset^r|}{|G_{[n]}^r|}$. Hence, no two tuples in R_B cover the same tuple in Q. Consequently, $\max_{B\in\mathcal{B}} |T_B^{\mathbf{D}}| \geq |R_{\bar{B}}|$, where $\bar{B} = \operatorname{argmin}_{B\in\mathcal{B}} |R_B|$. To complete the proof of (36), we are left to show $\log_2 |R_{\bar{B}}| \geq (1-\epsilon)\cdot k \cdot \min_{B\in\mathcal{B}} H[\mathbf{A}_B]$, which would follow immediately from showing that $\log_2 |R_{\bar{B}}| \geq (1-\epsilon)\cdot k \cdot H[\mathbf{A}_{\bar{B}}]$.

Recall that we assumed $\mathsf{LogSizeBound}_{\overline{\Gamma}_n^* \cap \mathsf{HDC}}(Q) \geq 1$ which means $H[\mathbf{A}_{\bar{B}}] \geq 1$, and thus $d \lg r \leq 2\delta r H[\mathbf{A}_{\bar{B}}]$ for large r. We have,

$$\begin{aligned}
\lg |R_{\bar{B}}| &= \log_2 \frac{r!}{|G_{\bar{B}}^r|} \\
&\geq \left[r\lg(r/e) + \frac{1}{2}\lg r + \lg\sqrt{2\pi} \right] \\
&\quad - \left[r\lg(r/e) - rH[\mathbf{A}_{\bar{B}}] + \sum_{\mathbf{a}_{\bar{B}}} \left(\lg e + \frac{1}{2}\lg(r\cdot p(\mathbf{a}_{\bar{B}})) \right) \right] \\
&= rH[\mathbf{A}_{\bar{B}}] + \lg\sqrt{2\pi} + \frac{1}{2}\lg r - \frac{1}{2}\sum_{\mathbf{a}_{\bar{B}}} \lg(re^2 p(\mathbf{a}_{\bar{B}})) \\
&\geq rH[\mathbf{A}_{\bar{B}}] + \frac{1}{2}\lg r - \frac{1}{2} d\lg(re^2) \\
&\geq r(1+\delta)H[\mathbf{A}_{\bar{B}}] - \delta r H[\mathbf{A}_{\bar{B}}] - \frac{1}{2} d\lg(r) + \frac{1}{2}(\lg r - d\lg(e^2)) \\
&\geq r(1+\delta)H[\mathbf{A}_{\bar{B}}] - \delta r H[\mathbf{A}_{\bar{B}}] - \frac{1}{2} d\lg(r) \\
&\geq r(1+\delta)H[\mathbf{A}_{\bar{B}}] - 2\delta r H[\mathbf{A}_{\bar{B}}] \\
&= \left(1 - \frac{2\delta}{1+\delta}\right) r(1+\delta) H[\mathbf{A}_{\bar{B}}]. \\
&= (1-\epsilon)\cdot k \cdot H[\mathbf{A}_{\bar{B}}].
\end{aligned}$$

□

4 Shannon flow inequalities and their proof sequences

The PANDA algorithm is built on the notion of a "proof sequence" for a specific class of Shannon inequalities of a particular form called the *Shannon flow inequalities*.

Definition 4.1 (Shannon flow inequality)**.** Let $\mathcal{B} \subseteq 2^V$ denote a collection of subsets of V. Let $\mathcal{C} \subseteq 2^V \times 2^V$ denote a collection of pairs (X, Y) such that $\emptyset \neq X \subset Y \subseteq V$. Let $\boldsymbol{\lambda}_{\mathcal{B}} = (\lambda_F)_{F\in\mathcal{B}} \in \mathbb{Q}_+^{\mathcal{B}}$ and $\boldsymbol{\delta}_{\mathcal{C}} = (\delta_{Y|X})_{(X,Y)\in\mathcal{C}} \in \mathbb{Q}_+^{\mathcal{C}}$ denote two vectors of non-negative rationals. For any polymatroid h, let $h(Y|X)$ denote $h(Y) - h(X)$.[13] If the following inequality

$$\sum_{B\in\mathcal{B}} \lambda_B \cdot h(B) \leq \sum_{(X,Y)\in\mathcal{C}} \delta_{Y|X} \cdot h(Y|X) \stackrel{\text{def}}{=} \sum_{(X,Y)\in\mathcal{C}} \delta_{Y|X} \cdot (h(Y) - h(X)) \tag{37}$$

[13]The quantity $h(Y|X)$ is the polymatroid-analog of the conditional entropy $H[Y|X] = H[Y] - H[X]$.

holds for all $h \in \Gamma_n$ (i.e. for all polymatroids), then it is called a *Shannon flow inequality*. The set $\mathcal{B}$ is called the set of *targets* of the flow inequality.

Section 4.1 motivates the study of Shannon flow inequalities. Section 4.2 explains why they are called "flow" inequalities. Then, Section 5 explains how the PANDA algorithm uses the results of this section.

4.1 Motivations for studying Shannon flow inequalities

Fix a disjunctive datalog query Q of the form (5) with degree constraints DC. Abusing notations, we write $(X, Y) \in \mathsf{DC}$ whenever $(X, Y, n_{Y|X}) \in \mathsf{DC}$. In particular the set DC can play the role of the generic set $\mathcal{C}$ in the definition of Shannon flow inequality. To explain where the Shannon flow inequalities come from, we study the (log) polymatroid output size bound (30) for disjunctive datalog query. Written more precisely (from (29)), the bound is

$$\mathsf{LogSizeBound}_{\Gamma_n \cap \mathsf{HDC}}(Q) \stackrel{\text{def}}{=} \max_{h \in \Gamma_n \cap \mathsf{HDC}} \min_{B \in \mathcal{B}} h(B).$$

We were able to replace sup by max because the set $\Gamma_n \cap \mathsf{HDC}$ is polyhedral. In particular, the above bound is the optimal objective value of the following optimization problem:

$$\begin{array}{lrll} \max & \min_{B \in \mathcal{B}} h(B) & & \\ \text{s.t.} & h(Y) - h(X) \le n_{Y|X} & (X, Y, n_{Y|X}) \in \mathsf{DC} & \text{(degree constraints)} \\ & h(I \cup J | J) - h(I | I \cap J) \le 0, & I \perp J & \text{(submodularity)} \\ & h(Y) - h(X) \ge 0, & \emptyset \neq X \subset Y \subseteq V & \text{(monotonicity)} \\ & h(Z) \ge 0, & \emptyset \neq Z \subseteq V. & \text{(non-negativity)} \end{array} \tag{38}$$

(Recall that implicitly we have $h(\emptyset) = 0$.) Here, $X \perp Y$ means $X \not\subseteq Y$ and $Y \not\subseteq X$. The optimization problem above is not easy to handle. We prove a simple lemma allowing us to reformulate it into a proper linear program. (A proof of the lemma is in Appendix B.)

Lemma 4.2. *Let $\mathbf{A} \in \mathbb{Q}^{\ell \times m}, \mathbf{b} \in \mathbb{R}^{\ell}$, and $\mathbf{C} \in \mathbb{Q}_+^{m \times p}$ be a matrix with columns $\mathbf{c}_1, \ldots, \mathbf{c}_p$. Consider the following maximin optimization problem:*

$$\begin{array}{ll} \max & \min_{k \in [p]} \mathbf{c}_k^T \mathbf{x} \\ s.t. & \mathbf{A}\mathbf{x} \le \mathbf{b} \\ & \mathbf{x} \ge \mathbf{0}. \end{array} \tag{39}$$

If problem (39)'s objective value is positive and bounded, then there exists a vector $\boldsymbol{\lambda} \in \mathbb{Q}_+^p$ satisfying the following conditions:

(a) $\|\boldsymbol{\lambda}\|_1 = 1$

(b) The problem (39) has the same optimal objective value as the following linear program:

$$\begin{array}{ll} \max & (\mathbf{C}\boldsymbol{\lambda})^T \mathbf{x} \\ s.t. & \mathbf{A}\mathbf{x} \le \mathbf{b} \\ & \mathbf{x} \ge \mathbf{0}. \end{array} \tag{40}$$

Back to reformulating the output size bound of disjunctive datalog, from Lemma 4.2, there exist non-negative coefficients $\boldsymbol{\lambda} = (\lambda_B)_{B \in \mathcal{B}}$ such that the optimization problem (38) has precisely the same objective

value as the following linear program:

$$\begin{aligned}
\max \quad & \sum_{B\in\mathcal{B}} \lambda_B \cdot h(B) && && (41)\\
\text{s.t.} \quad & h(Y)-h(X) \le n_{Y|X} && (X,Y,n_{Y|X})\in \mathsf{DC} && \text{(degree constraints)}\\
& h(I\cup J|J)-h(I|I\cap J)\le 0, && I\perp J && \text{(submodularity)}\\
& h(Y)-h(X)\ge 0, && \emptyset\neq X\subset Y\subseteq V && \text{(monotonicity)}\\
& h(Z)\ge 0, && \emptyset\neq Z\subseteq V. && \text{(non-negativity)}
\end{aligned}$$

It is from this linear program that Shannon flow inequalities arise, which is a simple consequence of Farkas lemma. The reasoning goes as follows. We first need the dual LP of (41). We associate a dual variable $\delta_{Y|X}$ to each degree constraint, a variable $\sigma_{I,J}$ to each submodularity constraint, and a variable $\mu_{X,Y}$ to each monotonicity constraint. For any $Z\in 2^V$, define the quantity

$$\mathsf{inflow}(Z) \stackrel{\text{def}}{=} \sum_{X:(X,Z)\in\mathsf{DC}} \delta_{Z|X} - \sum_{Y:(Z,Y)\in\mathsf{DC}} \delta_{Y|Z} + \sum_{\substack{I\perp J\\ I\cap J=Z}} \sigma_{I,J} + \sum_{\substack{I\perp J\\ I\cup J=Z}} \sigma_{I,J} - \sum_{J:J\perp Z} \sigma_{Z,J} - \sum_{X:X\subset Z} \mu_{X,Z} + \sum_{Y:Z\subset Y} \mu_{Z,Y}.$$

Note that the function $\mathsf{inflow}: 2^V\to\mathbb{Q}_+$ is also a function of the dual variables $(\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu})$. However, we do not explicitly write down this dependency to avoid heavy-loading notations. (Figure 3 illustrates the

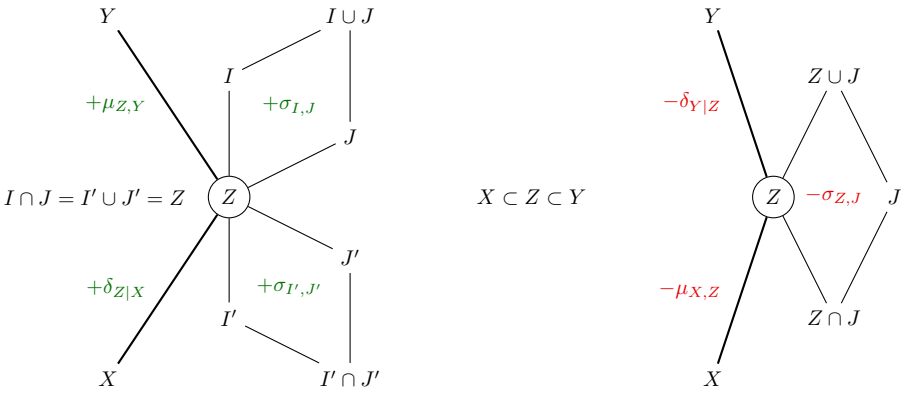

Figure 3: Contributions of coefficients to $\mathsf{inflow}(Z)$

contributions of various coefficients to $\mathsf{inflow}(Z)$ for a given set Z. It is instructive in the rest of the paper to refer back to this picture.) The dual of (41) is

$$\begin{aligned}
\min \quad & \textstyle\sum_{(X,Y)\in\mathsf{DC}} n_{Y|X}\cdot\delta_{Y|X} && && (42)\\
\text{s.t.} \quad & \mathsf{inflow}(B) && \ge \lambda_B,\ \forall B\in\mathcal{B}\\
& \mathsf{inflow}(Z) && \ge 0, \emptyset\neq Z\subseteq V.\\
& (\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu}) && \ge \mathbf{0}.
\end{aligned}$$

Let h^* denote an optimal solution to (41), then its objective value $h^*(V)$ is the same as the dual-optimal objective value due to strong duality. In particular, let $(\boldsymbol{\delta}^*,\boldsymbol{\sigma}^*,\boldsymbol{\mu}^*)$ denote an optimal solution to (42), then $\sum_{B\in\mathcal{B}}\lambda_B\cdot h^*(B) = \sum_{(X,Y)\in\mathsf{DC}}\delta^*_{Y|X}\cdot n_{Y|X}$; Instead of looking for h^* we can look for $(\boldsymbol{\delta}^*,\boldsymbol{\sigma}^*,\boldsymbol{\mu}^*)$ instead. In particular, to compute the bound we really only need the vector $\boldsymbol{\delta}^*$. One way to characterize $\boldsymbol{\delta}^*$, and in fact to characterize *any* dual feasible solution $(\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu})$, is to use Farkas' lemma [40], which in our context takes the following form (see Appendix B for a proof):

Proposition 4.3. *Given $\boldsymbol{\lambda}$ and $\boldsymbol{\delta}$, the inequality*

$$\sum_{B\in\mathcal{B}} \lambda_B \cdot h(B) \le \sum_{(X,Y)\in DC} \delta_{Y|X} \cdot h(Y|X) \tag{43}$$

is a Shannon flow inequality if and only if there exist $\boldsymbol{\sigma}$ and $\boldsymbol{\mu}$ such that $(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ is feasible to the dual LP (42).

In particular, the inequality holds when $\boldsymbol{\delta} = \boldsymbol{\delta}^*$, in which case the Shannon flow inequality implies the upper bound

$$\sum_{B\in\mathcal{B}} \lambda_B \cdot h(B) \le \sum_{(X,Y)\in \mathsf{DC}} \delta^*_{Y|X} \cdot n_{Y|X},$$

where the right hand side is the optimal objective value of both the primal and the dual.

4.2 Proof sequences for Shannon flow inequalities

A key observation from Abo Khamis et al. [32] was that we can turn a proof of a special case of inequality (43) into an algorithm. The proof has to be performed in a sequential manner; and this brings us to the concept of a proof sequence. In this paper, we refine the proof sequence notion from [32] in three significant ways. First, the definition of the proof sequence is different, allowing for a simpler algorithm (PANDA) than CSMA in [32]. Second, we are able to bound the length of the proof sequence to be polynomial in the size of the linear program (38), as opposed to the doubly exponential length in [32]. Third, new technical ideas are introduced so that we can construct proof sequences for the much more general Shannon flow inequality (43). The proof sequence from [32] was only for a special class of Shannon flow inequalities called "output inequality" in [32], captured by (68).

Definition 4.4 (Conditional polymatroids)**.** Let $\mathcal{P} \subseteq 2^V \times 2^V$ denote the set of all pairs (X, Y) such that $\emptyset \subseteq X \subset Y \subseteq V$. A vector $\mathbf{f} \in \mathbb{R}^{\mathcal{P}}_+$ has coordinates indexed by pairs $(X, Y) \in \mathcal{P}$. We denote the corresponding coordinate value of $\mathbf{f}$ by $f(Y|X)$. The vector $\mathbf{f}$ is called a *conditional polymatroid* iff there exists a polymatroid h such that $f(Y|X) = h(Y) - h(X)$; and, we say h defines the conditional polymatroid $\mathbf{f}$. Abusing notation somewhat, the conditional polymatroid defined by the polymatroid h is denoted by $\mathbf{h}$. In particular, $\mathbf{h} = (h(Y|X))_{(X,Y)\in\mathcal{P}}$. If h is a polymatroid then $h(\emptyset) = 0$, in which case we write $h(Y)$ instead of $h(Y|\emptyset)$.

To formally define the notion of a proof sequence, we rewrite the Shannon flow inequality (37) as an inequality on conditional polymatroids in the $\mathbb{Q}^{\mathcal{P}}_+$ space. We extend the vectors $\boldsymbol{\lambda}_{\mathcal{B}} \in \mathbb{Q}^{\mathcal{B}}_+$ and $\boldsymbol{\delta}_{\mathcal{C}} \in \mathbb{Q}^{\mathcal{C}}_+$ to become vectors $\boldsymbol{\lambda}, \boldsymbol{\delta}$ in the $\mathbb{Q}^{\mathcal{P}}_+$ space in the obvious way:

$$\lambda(Y|X) \stackrel{\text{def}}{=} \begin{cases} \lambda_{\mathcal{B}}(B) & \text{when } Y = B, X = \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

$$\delta(Y|X) \stackrel{\text{def}}{=} \begin{cases} \delta_{\mathcal{C}}(Y|X) & \text{when } (X, Y) \in \mathcal{C} \\ 0 & \text{otherwise.} \end{cases}$$

Then, inequality (37) can be written simply as $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$. Note the crucial fact that, even though $\boldsymbol{\lambda} \in \mathbb{Q}^{\mathcal{P}}_+$, for it to be part of a Shannon flow inequality only the entries $\lambda_{B|\emptyset}$ can be positive. We will often write λ_B instead of $\lambda_{B|\emptyset}$. These assumptions are implicit henceforth. Proposition 4.3 can now be written simply as:

Proposition 4.5. *Given any $\boldsymbol{\lambda}, \boldsymbol{\delta} \in \mathbb{Q}^{\mathcal{P}}_+$, where $\lambda_{Y|X} > 0$ implies $X = \emptyset$, the inequality $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ is a Shannon flow inequality if and only if there exist $\boldsymbol{\sigma}$ and $\boldsymbol{\mu}$ such that $(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ belongs to the polyhedron*

$$\big\{(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu}) \mid \mathsf{inflow}(Z) \ge \lambda_Z, \emptyset \ne Z \subseteq V, (\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu}) \ge \mathbf{0}\big\}. \tag{44}$$

The conditional polymatroids satisfy four basic laws:

$$\begin{aligned} h(I\cup J|J) - h(I|I\cap J) &\le 0, \quad I \perp J \text{ (submodularity)} \\ -h(Y|\emptyset) + h(X|\emptyset) &\le 0, \quad X \subset Y \text{ (monotonicity)} \\ h(Y|\emptyset) - h(Y|X) - h(X|\emptyset) &\le 0, \quad X \subset Y \text{ (composition)} \\ -h(Y|\emptyset) + h(Y|X) + h(X|\emptyset) &\le 0, \quad X \subset Y \text{ (decomposition)} \end{aligned}$$

(We will explain later why the last two relations are named composition and decomposition laws.) For every $I \perp J$, define a vector $\mathbf{s}_{I,J} \in \mathbb{Q}_+^{\mathcal{P}}$, and for every $X \subset Y$, define three vectors $\mathbf{m}_{X,Y}, \mathbf{c}_{X,Y}, \mathbf{d}_{Y,X} \in \mathbb{Q}_+^{\mathcal{P}}$ such that the laws above can be written correspondingly in dot-product form:

$$\begin{aligned} \langle \mathbf{s}_{I,J}, \mathbf{h}\rangle &\le 0, \quad I \perp J \text{ (submodularity)} && (45)\\ \langle \mathbf{m}_{X,Y}, \mathbf{h}\rangle &\le 0, \quad X \subset Y \text{ (monotonicity)} && (46)\\ \langle \mathbf{c}_{X,Y}, \mathbf{h}\rangle &\le 0, \quad X \subset Y \text{ (composition)} && (47)\\ \langle \mathbf{d}_{Y,X}, \mathbf{h}\rangle &\le 0, \quad X \subset Y \text{ (decomposition)} && (48) \end{aligned}$$

Definition 4.6 (Inequality sequence)**.** An *inequality sequence* is a sequence of *non-negative* vectors $\boldsymbol{\delta}_0, \boldsymbol{\delta}_1, \ldots, \boldsymbol{\delta}_\ell$ such that, for every $i \in [\ell]$, $\boldsymbol{\delta}_i = \boldsymbol{\delta}_{i-1} + w_i \cdot \mathbf{f}_i$ for some $w_i > 0$ and $\mathbf{f}_i \in \{\mathbf{s}_{I,J}, \mathbf{m}_{X,Y}, \mathbf{c}_{X,Y}, \mathbf{d}_{Y,X}\}$. The reason this is called an "inequality" sequence is that, obviously, $\langle \boldsymbol{\delta}_{i-1}, \mathbf{h}\rangle \ge \langle \boldsymbol{\delta}_i, \mathbf{h}\rangle$ for every conditional polymatroid $\mathbf{h}$, and every $i \in [\ell]$.

Definition 4.7 (Proof sequence)**.** Given a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$. Suppose there is an inequality sequence $\boldsymbol{\delta}_0, \boldsymbol{\delta}_1, \ldots, \boldsymbol{\delta}_\ell$ such that $\boldsymbol{\delta}_0 = \boldsymbol{\delta}$, and $\boldsymbol{\delta}_\ell \ge \boldsymbol{\lambda}$; then the sequence $(w_1\mathbf{f}_1, \ldots, w_\ell\mathbf{f}_\ell)$ is called the *proof sequence* of the inequality, where w_i is the *weight* of the *proof step* $\mathbf{f}_i$. The step $\mathbf{s}_{I,J}$ is called a *submodularity step*, $\mathbf{m}_{X,Y}$ a *monotonicity step*, $\mathbf{d}_{Y,X}$ a *decomposition step*, and $\mathbf{c}_{X,Y}$ a *composition step*.

(Note that our notion of a proof sequence here is *different* from that in [32]. This notion of a proof sequence helps make PANDA a cleaner/simpler algorithm, and also helps reduce the bound on the length of a good proof sequence.) Before we prove the existence of a proof sequence for the general Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$, we introduce some notions.

Definition 4.8 (Witness)**.** Let $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ be a Shannon flow inequality. From Proposition 4.3 there exists $(\boldsymbol{\sigma}, \boldsymbol{\mu})$ such that $(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ belongs to the polyhedron (44). We call $(\boldsymbol{\sigma}, \boldsymbol{\mu})$ a *witness* for the Shannon flow inequality. A witness is said to be *tight* if $\mathsf{inflow}(Z) = \lambda_Z$ for all Z.

We remark that, if $\mathsf{inflow}(Z) > \lambda_Z$, we can always increase $\mu_{\emptyset,Z}$ by the amount $\mathsf{inflow}(Z) - \lambda_Z$ so that $\mathsf{inflow}(Z) = \lambda_Z$. In particular, it is easy to turn any witness into a tight witness. We provide below a simple construction of a proof sequence for a given Shannon flow inequality. (Appendix B presents more advanced constructions with smaller bounds on the length of the proof sequence.)

Theorem 4.9 (Constructing a proof sequence)**.** *Let $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ be a Shannon flow inequality with witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$. There exists a proof sequence for the inequality $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ with length at most $D(3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1)$, where D is the minimum common denominator of all entries in $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$.*

Proof. We prove this fact by induction on the quantity $\ell(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu}) \stackrel{\text{def}}{=} D\left(\|\boldsymbol{\lambda}\|_1 + 2\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1\right)$, which is an integer. The base case is when $\|\boldsymbol{\lambda}\|_1 = 0$, which is trivial because the inequality has a proof sequence of length 0. In the inductive step, assume $\|\boldsymbol{\lambda}\|_1 > 0$, in which case there must be some $B \subseteq V$ for which $\lambda_B > 0$. Our plan is to produce a Shannon flow inequality $\langle \boldsymbol{\lambda}', \mathbf{h}\rangle \le \langle \boldsymbol{\delta}', \mathbf{h}\rangle$ witnessed by $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ such that $\ell(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}') < \ell(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$. Then, from the induction hypothesis we obtain a proof sequence $\mathsf{ProofSeq}'$ for $\langle \boldsymbol{\lambda}', \mathbf{h}\rangle \le \langle \boldsymbol{\delta}', \mathbf{h}\rangle$. Finally the proof sequence $\mathsf{ProofSeq}$ for $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ is constructed from $\mathsf{ProofSeq}'$ by appending to the beginning one or two proof steps.

More concretely, from Proposition 4.5, we know $\sum_{\emptyset \ne W \subseteq V} \mathsf{inflow}(W) \ge \lambda_B > 0$. Consequently, there must exist some $Z \ne \emptyset$ for which $\delta_{Z|\emptyset} > 0$; otherwise, all dual variables $\delta_{Y|X}, \sigma_{I,J}, \mu_{X,Y}$ contribute a non-positive amount to the sum $\sum_{\emptyset \ne W \subseteq V} \mathsf{inflow}(W)$. Let $w \stackrel{\text{def}}{=} 1/D$, and fix an arbitrary $Z \ne \emptyset$ where $\delta_{Z|\emptyset} > 0$.

We initially set $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}') = (\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$; then we modify $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$ slightly depending on the cases below.

Case 1: $\lambda_Z > 0$. Reduce both λ'_Z and $\delta'_{Z|\emptyset}$ by w. From Proposition 4.5, we can verify that $\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ is a Shannon flow inequality witnessed by $(\boldsymbol{\sigma}', \boldsymbol{\mu}') = (\boldsymbol{\sigma}, \boldsymbol{\mu})$. By induction hypothesis, $\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ has a proof sequence $\mathsf{ProofSeq}'$ of length at most $D(3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}'\|_1 + \|\boldsymbol{\mu}\|_1)$. Furthermore, the $\mathsf{ProofSeq}'$ *is* also a proof sequence for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$.

Case 2: $\lambda_Z = 0$ *and* $\mathsf{inflow}(Z) > 0$. Reduce $\delta'_{Z|\emptyset}$ by w. Then, from Proposition 4.5, we can verify that $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ is a Shannon flow inequality witnessed by $(\boldsymbol{\sigma}, \boldsymbol{\mu})$. The inductive step is now identical to that of Case 1.

Case 3: $\mathsf{inflow}(Z) = 0$. Since $\delta_{Z|\emptyset} > 0$, there must be some dual variable that is contributing a negative amount to $\mathsf{inflow}(Z)$. (See the right side of Figure 3.) In particular, one of the following three cases must hold

- There is some $X \subset Z$ such that $\mu_{X,Z} \ge w$. Define $\boldsymbol{\delta}' = \boldsymbol{\delta} + w \cdot \mathbf{m}_{X,Z}$ and reduce $\mu'_{X,Z}$ by w. Note that $\|\boldsymbol{\delta}'\|_1 = \|\boldsymbol{\delta}\|_1$, $\|\boldsymbol{\mu}'\|_1 = \|\boldsymbol{\mu}\|_1 - w$, and $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ is a Shannon flow inequality. witnessed by $(\boldsymbol{\sigma}, \boldsymbol{\mu}')$. By induction hypothesis, $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ has a proof sequence $\mathsf{ProofSeq}'$ of length at most $D(3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}'\|_1 + \|\boldsymbol{\mu}'\|_1)$. It follows that $\mathsf{ProofSeq} = (w \cdot \mathbf{m}_{X,Z}, \mathsf{ProofSeq}')$ is a proof sequence for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ of length at most $D(3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1)$.
- There is some $Y \supset Z$ such that $\delta_{Y|Z} \ge w$. Define $\boldsymbol{\delta}' = \boldsymbol{\delta} + w \cdot \mathbf{c}_{Z,Y}$. Note that $\|\boldsymbol{\delta}'\|_1 = \|\boldsymbol{\delta}\|_1 - w$ and $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ is a Shannon flow inequality witnessed by $(\boldsymbol{\sigma}, \boldsymbol{\mu})$. From the proof sequence $\mathsf{ProofSeq}'$ for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ we obtain the proof sequence $\mathsf{ProofSeq} = (w \cdot \mathbf{c}_{Z,Y}, \mathsf{ProofSeq}')$ for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ of the desired length.
- There is some $J \perp Z$ such that $\sigma_{Z,J} \ge w$. Define $\boldsymbol{\delta}' = \boldsymbol{\delta} + w \cdot \mathbf{d}_{Z,Z\cap J} + w \cdot \mathbf{s}_{Z,J}$, and reduce $\sigma'_{Z,J}$ by w. In this case, $\|\boldsymbol{\delta}'\|_1 = \|\boldsymbol{\delta}\|_1 + w$, $\|\boldsymbol{\sigma}'\|_1 = \|\boldsymbol{\sigma}\|_1 - w$, and $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ is a Shannon flow inequality witnessed by $(\boldsymbol{\sigma}', \boldsymbol{\mu})$. By induction hypothesis, $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ has a proof sequence $\mathsf{ProofSeq}'$ of length at most $D(3\|\boldsymbol{\sigma}'\|_1 + \|\boldsymbol{\delta}'\|_1 + \|\boldsymbol{\mu}\|_1)$. It follows that $\mathsf{ProofSeq} = (w \cdot \mathbf{d}_{Z,Z\cap J}, w \cdot \mathbf{s}_{Z,J}, \mathsf{ProofSeq}')$ is a proof sequence for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ of length at most $D(3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1)$.

□

The PANDA algorithm needs another technical lemma regarding Shannon flow inequalities.

Lemma 4.10 (Truncating a Shannon flow inequality)**.** *Let $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ be a Shannon flow inequality with witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$. Let D be a common denominator of all entries in $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$, and $w = 1/D$. Suppose $\|\boldsymbol{\lambda}\|_1 > 0$ and $\delta_{Y|\emptyset} > 0$. Then, there are two vectors $\boldsymbol{\lambda}', \boldsymbol{\delta}'$ satisfying the following conditions:*

(a) $\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ is a Shannon flow inequality (with witness $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$).

(b) $\boldsymbol{\lambda}' \le \boldsymbol{\lambda}$ and $\boldsymbol{\delta}' \le \boldsymbol{\delta}$ (component-wise comparisons).

(c) $\|\boldsymbol{\lambda}'\|_1 \ge \|\boldsymbol{\lambda}\|_1 - w$ and $\delta'_{Y|\emptyset} \le \delta_{Y|\emptyset} - w$.

(d) D is a common denominator for all entries in $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$.

The proof follows the same strategy as in the proof of Theorem 4.9, but with some subtle differences.

Proof. W.L.O.G. we can assume that $(\boldsymbol{\sigma}, \boldsymbol{\mu})$ is a tight witness. We construct $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$ from $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ as follows. Initially we set $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}') = (\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$. Let $\mathsf{inflow}'(Z)$ denote the quantity $\mathsf{inflow}(Z)$ measured on the vector $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$. Due to tightness of the witness, at this point $\mathsf{inflow}'(Z) - \lambda'_Z = 0$ for every Z.

Now, we start disturbing the flow balance equations starting from $Z = Y$ by setting $\delta'_{Z|\emptyset} = \delta'_{Z|\emptyset} - w$ which means $\mathsf{inflow}'(Z)$ was reduced by w. Note that Z is the only point for which $\mathsf{inflow}'(Z) - \lambda'_Z \neq 0$ (it is

negative). If $\lambda'_Z > 0$, then we simply reduce λ'_Z by w and terminate. If $\lambda'_Z = 0$, then either (1) there is some $X \subset Z$ such that $\mu'_{X,Z} \geq w$, (2) there is some $Y \supset Z$ such that $\delta'_{Y|Z} \geq w$, or (3) there is some $J \perp Z$ such that $\sigma'_{Z,J} \geq w$. (Check Figure 3.) Cases (1) and (2) are handled in a similar way to the proof of Theorem 4.9 while case (3) is handled differently. In particular, if (1) holds, then we reduce $\mu'_{X,Z}$ by w and set $Z = X$. If (2) holds, then we reduce $\delta'_{Y|Z}$ by w and set $Z = Y$. If (3) holds, then we reduce $\sigma'_{Z,J}$ by w, *increase* $\mu'_{Z\cap J,J}$ by w, and set $Z = Z \cup J$. In all three cases, (the new) Z is the only point where $\mathsf{inflow}'(Z) - \lambda'_Z$ has a deficit, and the process continues if the new $Z \neq \emptyset$.

The above process terminates because every time we move Z to a new deficit point, the quantity $2\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1$ is reduced by w. When the process terminates, all quantities $\mathsf{inflow}'(Z) - \lambda'(Z) = 0$ and thus $\langle \boldsymbol{\lambda}', \mathbf{h}\rangle \leq \langle \boldsymbol{\delta}', \mathbf{h}\rangle$ is a Shannon flow inequality witnessed by $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ by Proposition 4.5. Property (b) holds because we only reduce the $\boldsymbol{\sigma}'$ and $\boldsymbol{\delta}'$ entries. Properties (c) and (d) hold trivially. □

Appendix 3 presents bounds on $\|\boldsymbol{\sigma}\|_1$, $\|\boldsymbol{\delta}\|_1$ and $\|\boldsymbol{\mu}\|_1$ for a given Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \leq \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ with witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$. By plugging those bounds into Theorem 4.9, we can bound the length of the constructed proof sequence by $D \cdot (\frac{3}{2}n^3 + n)$. (See Corollary B.7.) Appendix B.3 presents an advanced construction of a proof sequence whose length is polynomial in the size of the linear program (38). The construction relies on connections to flow networks and uses Edmond-Karp algorithm [18] as a black-box.

5 The PANDA algorithm

This section presents an algorithm called PANDA, to evaluate a disjunctive datalog query Q in time predicted by its polymatroid bound (see Definition 29):

$$\tilde{O}(N + \mathrm{poly}(\log N) \cdot 2^{\mathsf{LogSizeBound}_{\Gamma_n \cap \mathsf{HDC}}(Q)}).$$

In particular, the algorithm would be worst-case optimal (up to poly-log factor) if $\mathsf{LogSizeBound}_{\Gamma_n\cap\mathsf{HDC}}(Q)$ is tight. We knew from Section 3 that $\mathsf{LogSizeBound}_{\Gamma_n\cap\mathsf{HDC}}$ is not tight in general. However, there are known classes of queries with FDs for which this bound is tight [32]. To be more formal, we define the disjunctive datalog query evaluation problem below. Section 6.3 explains how PANDA is used to achieve the degree-aware fractional hypertree width and degree-aware submodular width.

Definition 5.1 (DDQE)**.** An instance of the *disjunctive datalog query evaluation* problem (DDQE) is comprised of a collection $\mathcal{R}$ of input relations R_F, for $F \in \mathcal{E}$, where $\mathcal{H} = (V, \mathcal{E})$ is the query's hypergraph, and a set of degree constraints DC containing triples $(X, Y, n_{Y|X})$, where $X \subset Y \subseteq V$, and $n_{Y|X}$ is a non-negative real number. As usual, each degree constraint has an input relation guarding it (see Definition 2.8). In addition, there is a collection $\mathcal{B} \subseteq 2^V \setminus \{\emptyset\}$ of "targets" for the problem. Let Q denote the join of all input relations (with degree constraints DC). The output of the DDQE problem is a tuple $\mathbf{T} = (T_B)_{B\in\mathcal{B}}$ of relations, where relation T_B is on attributes $\mathbf{A}_B$, and for every tuple $\mathbf{a} \in Q$, there exists $B \in \mathcal{B}$ for which $\Pi_B(\mathbf{a}) \in T_B$.

The algorithm sketch is shown in the box Algorithm 1. It takes as input the collection of input relations $\mathcal{R}$, the degree constraints DC, a Shannon flow inequality and its proof sequence. The Shannon flow inequality is constructed by solving the optimization problem $\max_{h\in\Gamma_n\cap\mathsf{HDC}} \min_{B\in\mathcal{B}} h(B)$. From Lemma 4.2, we can find a vector $\boldsymbol{\lambda}_{\mathcal{B}}$ with $\|\boldsymbol{\lambda}\|_1 = 1$ such that the problem has the same optimal objective value as the linear program $\max_{h\in\Gamma_n\cap\mathsf{HDC}}\langle \boldsymbol{\lambda}, \mathbf{h}\rangle$. Recall from Section 4.2 that, when we extend $\boldsymbol{\lambda}_{\mathcal{B}}$ to the (conditional polymatroid) space $\boldsymbol{\lambda} \in \mathbb{Q}_+^{\mathcal{P}}$, only the entries $\lambda_{B|\emptyset}$ for $B \in \mathcal{B}$ can be positive. Let $(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ denote an optimal dual solution to this LP, then from Proposition 4.3 we know $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \leq \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ is a Shannon flow inequality. For inductive purposes, we will assume a slightly more general condition that the guiding Shannon flow inequality satisfies $0 < \|\boldsymbol{\lambda}\|_1 \leq 1$. From Theorem 4.9, we obtain a proof sequence for the Shannon flow inequality. Also for inductive purposes, the following invariant is maintained throughout:

Definition 5.2 (Degree-support invariant)**.** The guiding Shannon flow inequality is promised to satisfy the following invariant: if $\delta_{Y|X} > 0$ then there exist $Z \subseteq X$, $W \subseteq Y$ such that $W - Z = Y - X$ and

Algorithm 1 PANDA($\mathcal{R}$, DC, $(\boldsymbol{\lambda}, \boldsymbol{\delta})$, ProofSeq)

Input: $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ is a Shannon flow inequality with proof sequence ProofSeq, $0 < \|\boldsymbol{\lambda}\|_1 \le 1$
Input: DC are input degree constraints guarded by input relations $\mathcal{R}$
Input: The degree-support invariant (Definition 5.2) is satisfied
Input: $N \le 2^{\mathsf{OBJ}}$, where N is the input size and OBJ is given by (49)

If an input relation $R \in \mathcal{R}$ has attribute set $\mathbf{A}_B$, with $B \in \mathcal{B}$ **then**
  Return $T_B = R$ ▷ Only one table in the output
Let ProofSeq $= (w \cdot \mathbf{f}, \mathsf{ProofSeq}')$ ▷ $\mathbf{f}$ is the first proof step, with weight w
$\boldsymbol{\delta}' \leftarrow \boldsymbol{\delta} + w \cdot \mathbf{f}$ ▷ Advance the proof step
If $\mathbf{f} = \mathbf{s}_{I,J}$ **then** ▷ Case 1. $\delta_{I|I\cap J} \ge w > 0$ must hold
  Return PANDA($\mathcal{R}$, DC, $(\boldsymbol{\lambda}, \boldsymbol{\delta}')$, ProofSeq$'$)
else If $\mathbf{f} = \mathbf{m}_{X,Y}$ **then** ▷ Case 2. $\delta_{Y|\emptyset} \ge w > 0$ must hold
  Let R be a guard for $(\emptyset, Y, n_{Y|\emptyset}) \in \mathsf{DC}$, which supports $\delta_{Y|\emptyset}$
  $\mathcal{R}' = \mathcal{R} \cup \{\Pi_X(R)\}$
  $\mathsf{DC}' = \mathsf{DC} \cup \{(\emptyset, X, \log_2 |\Pi_X(R)|)\}$
  Return PANDA($\mathcal{R}'$, DC$'$, $(\boldsymbol{\lambda}, \boldsymbol{\delta}')$, ProofSeq$'$)
else If $\mathbf{f} = \mathbf{d}_{Y,X}$ **then** ▷ Case 3. $\delta_{Y|\emptyset} \ge w > 0$ must hold
  Let $R \in \mathcal{R}$ be a guard for $(\emptyset, Y, n_{Y|\emptyset}) \in \mathsf{DC}$
  Partition $R = R^{(1)} \cup \cdots \cup R^{(k)}$ as in (51) ▷ $O(\log_2 |R|)$ branches
  For $j \leftarrow 1$ **to** k **do**
    $\mathcal{R}^{(j)} \leftarrow \mathcal{R} \cup \{R^{(j)}\}$
    $\mathsf{DC}^{(j)} = \mathsf{DC} \cup \{(\emptyset, X, n^{(j)}_{X|\emptyset}), (X, Y, n^{(j)}_{Y|X})\}$
    $(T^{(j)}_B)_{B\in\mathcal{B}} \leftarrow$ PANDA$(\mathcal{R}^{(j)}, \mathsf{DC}^{(j)}, (\boldsymbol{\lambda}, \boldsymbol{\delta}'), \mathsf{ProofSeq}')$
  Return $\left(T_B \stackrel{\text{def}}{=} \bigcup_{j=1}^{k} T^{(j)}_B\right)_{B\in\mathcal{B}}$ ▷ Union of results from branches
else If $\mathbf{f} = \mathbf{c}_{X,Y}$ **then** ▷ Case 4. $\delta_{Y|X} \ge w > 0$ and $\delta_{X|\emptyset} \ge w > 0$
  Let R be a guard of $(Z, W, n_{W|Z}) \in \mathsf{DC}$, which supports $\delta_{Y|X}$
  Let S be a guard of $(\emptyset, X, n_{X|\emptyset}) \in \mathsf{DC}$, which supports $\delta_{X|\emptyset}$
  If $n_{W|Z} + n_{X|\emptyset} \le \mathsf{OBJ}$ **then** ▷ Case 4a. OBJ defined in (49)
    Compute $T(\mathbf{A}_Y) \leftarrow \Pi_X(S) \bowtie \Pi_W(R)$ ▷ With runtime $\tilde{O}(2^{\mathsf{OBJ}})$
    $\mathcal{R}' = \mathcal{R} \cup \{T\}$
    $\mathsf{DC}' = \mathsf{DC} \cup \{(\emptyset, Y, \log_2 |T|)\}$
    Return PANDA($\mathcal{R}'$, DC$'$, $(\boldsymbol{\lambda}, \boldsymbol{\delta}')$, ProofSeq$'$)
  else ▷ Case 4b.
    Let $\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ be the truncated Shannon flow inequality stated in Lemma 4.10
    Recompute a fresh ProofSeq$'$ for $\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$
    Return PANDA($\mathcal{R}$, DC, $(\boldsymbol{\lambda}', \boldsymbol{\delta}')$, ProofSeq$'$)

$(Z, W, n_{W|Z}) \in \mathsf{DC}$. (Note that, if $X = \emptyset$ then $W = Y$ and $Z = \emptyset$.) The degree constraint $(Z, W, n_{W|Z})$ is said to *support* the positive $\delta_{Y|X}$. (If there are multiple constraints $(Z, W, n_{W|Z})$ supporting $\delta_{Y|X}$, then the the one with the minimum $n_{W|Z}$ is always chosen to be *the* constraint that supports $\delta_{Y|X}$, where ties are broken arbitrarily.)

Invariant 5.2 is certainly satisfied in the very beginning. A very high-level description of the algorithm is as follows. Recall from Definition 4.7 that a proof sequence ProofSeq is a series of proof steps $w \cdot \mathbf{f}$: submodularity ($\mathbf{f} = \mathbf{s}_{I,J}$), monotonicity ($\mathbf{f} = \mathbf{m}_{X,Y}$), decomposition ($\mathbf{f} = \mathbf{d}_{Y,X}$), and composition ($\mathbf{f} = \mathbf{c}_{X,Y}$). (The vectors $\mathbf{s}_{I,J}$, $\mathbf{m}_{X,Y}$, $\mathbf{d}_{Y,X}$, and $\mathbf{c}_{X,Y}$ were defined in equations (45)– (48).) PANDA takes the proof sequence as "symbolic instructions". For each instruction, PANDA does some computation, spawns a number of subproblem(s) all of which are instances of DDQE, and creates new (intermediate) relations to become

input of the subproblems if necessary. The output of the ith subproblem is a set of tables $T_B^{(i)}$ for $B \in \mathcal{B}$. The overall output is the set of tables $T_B = \bigcup_i T_B^{(i)}$, $B \in \mathcal{B}$; namely for each $B \in \mathcal{B}$ we take the union of the corresponding tables from the subproblems's outputs. The number of subproblems will be shown to be polylogarithmic in the input size.

We next walk the reader step-by-step through the algorithm. The algorithm stops as soon as there is a relation $R \in \mathcal{R}$ with attribute set $\mathbf{A}_B$ where $B \in \mathcal{B}$, in which case R is the output. Define the quantity

$$n(\delta_{Y|X}) = \begin{cases} \delta_{Y|X} \cdot n_{W|Z} & \text{if } \delta_{Y|X} > 0 \text{ and } (Z, W, n_{W|X}) \text{ supports it} \\ 0 & \text{if } \delta_{Y|X} = 0, \end{cases}$$

and define a "budget" quantity of

$$\mathsf{OBJ} \stackrel{\text{def}}{=} \sum_{(X,Y)} n(\delta_{Y|X}). \tag{49}$$

We assume that the input size is within the budges, i.e.

$$N \le 2^{\mathsf{OBJ}}. \tag{50}$$

(We will see later how to enforce this assumption when it is not satisfied.) We will keep every step of the algorithm to run within the budget of $\tilde{O}(2^{\mathsf{OBJ}})$. Specifically, we will keep every table that the algorithm computes of size $\le 2^{\mathsf{OBJ}}$. The algorithmic steps are modeled after the proof steps. Let $\mathbf{f}$ be the first proof step (instruction) with weight w, i.e. $\mathsf{ProofSeq} = (w \cdot \mathbf{f}, \mathsf{ProofSeq}')$ where $\mathsf{ProofSeq}'$ contains the rest of the instructions.

Case 1: $\mathbf{f} = \mathbf{s}_{I,J}$ is a submodularity step. By definition of proof sequence, $\boldsymbol{\delta} + w \cdot \mathbf{s}_{I,J} \ge \mathbf{0}$, and thus $\delta_{I|I\cap J} \ge w > 0$. Let $(Z, W, n_{W|Z}) \in \mathsf{DC}$ be the degree constraint supporting $\delta_{I|I\cap J}$; then $Z \subseteq I \cap J$, $W \subseteq I$, and $W - Z = I - I \cap J$. The algorithm proceeds by setting $\boldsymbol{\delta}' = \boldsymbol{\delta} + w \cdot \mathbf{s}_{I,J}$. Note that $\delta'_{I\cup J|J}$ is now positive, and so it needs a supporting degree constraint. From the fact that $W - Z = I - I \cap J = I \cup J - J$, $(Z, W, n_{W|Z})$ can support $\delta'_{I\cup J|J}$. Since $\mathbf{f}$ was the next step in the proof sequence, $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}', \mathbf{h}\rangle$ is a Shannon flow inequality with proof sequence $\mathsf{ProofSeq}'$.

Case 2: $\mathbf{f} = \mathbf{m}_{X,Y}$ is a monotonicity step. By definition of proof sequence, $\boldsymbol{\delta} + w \cdot \mathbf{m}_{X,Y} \ge \mathbf{0}$, and thus $\delta_{Y|\emptyset} \ge w > 0$. Let $(\emptyset, Y, n_{Y|\emptyset}) \in \mathsf{DC}$ be the degree constraint supporting $\delta_{Y|\emptyset}$, and R be an input relation guarding the degree constraint (hence by Definition 2.8, $\log_2 |\Pi_Y(R)| \le n_{Y|\emptyset}$). Then, we proceed by setting $\boldsymbol{\delta}' = \boldsymbol{\delta} + w \cdot \mathbf{m}_{X,Y}$. Note that $\delta'_{X|\emptyset}$ is positive, and so it needs a supporting degree constraint, which is the newly added degree constraint $(\emptyset, X, n_{X|\emptyset})$, guarded by $\Pi_X(R)$, where $n_{X|\emptyset} \stackrel{\text{def}}{=} \log_2 |\Pi_X(R)| \le \log_2 |R| \le \mathsf{OBJ}$.

Case 3: $\mathbf{f} = \mathbf{d}_{Y,X}$ is a decomposition step with weight w. From $\boldsymbol{\delta} + w \cdot \mathbf{d}_{Y,X} \ge \mathbf{0}$, it follows that $\delta_{Y|\emptyset} \ge w > 0$. From the guarantee that $\delta_{Y|\emptyset}$ has a supporting degree constraint, it follows that there is a relation $R \in \mathcal{R}$ guarding $(\emptyset, Y, n_{Y|\emptyset})$, which means $\log_2 |\Pi_Y(R)| \le n_{Y|\emptyset}$. Lemma 5.3 shows that R can be partitioned into at most $(k = 2\log_2 |R| \le 2 \cdot \mathsf{OBJ})$ sub-tables $R^{(1)}, \ldots, R^{(k)}$ such that $n^{(j)}_{Y|X} + n^{(j)}_{X|\emptyset} \le n_{Y|\emptyset}$, for all $j \in [k]$, where

$$\begin{aligned} n^{(j)}_{Y|X} &\stackrel{\text{def}}{=} \max_{\mathbf{t} \in \Pi_X(R^{(j)})} \log_2 \deg_{R^{(j)},Y}(\mathbf{t}), \\ n^{(j)}_{X|\emptyset} &\stackrel{\text{def}}{=} \log_2 |\Pi_X(R^{(j)})|. \end{aligned} \tag{51}$$

(See (25) in Definition 2.8.) For each of these sub-tables $R^{(j)}$ of R, we create a subproblem with the same input tables but with R replaced by $R^{(j)}$. The jth subproblem has degree constraints $\mathsf{DC}^{(j)}$ where

$$\mathsf{DC}^{(j)} = \mathsf{DC} \cup \big\{(\emptyset, X, n^{(j)}_{X|\emptyset}), (X, Y, n^{(j)}_{Y|X})\big\}.$$

The table $R^{(j)}$ guards both of the new degree constraints. Set $\boldsymbol{\delta} = \boldsymbol{\delta} + w \cdot \mathbf{d}_{Y,X}$. The jth subproblem is on the Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$ with proof sequence $\mathsf{ProofSeq}'$.

Case 4: $\mathbf{f} = \mathbf{c}_{X,Y}$ is a composition step with weight w. By definition of proof sequence, $\boldsymbol{\delta} + w \cdot \mathbf{c}_{X,Y} \geq \mathbf{0}$, and thus $\delta_{Y|X} \geq w > 0$ and $\delta_{X|\emptyset} \geq w > 0$. From the facts that $\delta_{Y|X}$ and $\delta_{X|\emptyset}$ have supports, there must be two sets $Z \subseteq X$ and $W \subseteq Y$ for which $W - Z = Y - X$ and $(Z, W, n_{W|Z}) \in \mathsf{DC}$ which is guarded by an input relation R; and an input relation S for which $\log_2 |\Pi_X(S)| \leq n_{X|\emptyset}$. Note that $X \cup (W - Z) = X \cup (Y - X) = Y$. We consider two cases:

(Case 4a) If $n_{W|Z} + n_{X|\emptyset} \leq \mathsf{OBJ}$, then we can compute the table $T(\mathbf{A}_Y) \stackrel{\text{def}}{=} \Pi_X(S) \bowtie \Pi_W(R)$ by going over all tuples in $\Pi_X(S)$ and expanding them using matching tuples in $\Pi_W(R)$. The runtime of the join is $\tilde{O}(2^{n_{X|\emptyset} + n_{W|Z}}) = \tilde{O}(2^{\mathsf{OBJ}})$, and the size of T is $\leq 2^{\mathsf{OBJ}}$. As for the time it takes to compute the projections $\Pi_X(S), \Pi_W(R)$, we reason as follows. If S and/or R was an input relation, then the cost of compute these projections can be amortized into the term $\tilde{O}(N)$ of the runtime. If S and/or R was an intermediate relation, then their sizes cannot exceed 2^{OBJ} and the cost of computing projections can be amortized into the term $\tilde{O}(2^{\mathsf{OBJ}})$ in the runtime. The Shannon flow inequality is modified by setting $\boldsymbol{\delta} = \boldsymbol{\delta} + w \cdot \mathbf{c}_{X,Y}$, with the the proof sequence $\mathsf{ProofSeq}'$, and the set of degree constraints is extended by adding the constraint $(\emptyset, Y, n_{Y|\emptyset} \stackrel{\text{def}}{=} \log_2 |T|))$, guarded by T.

(Case 4b) If $n_{W|Z} + n_{X|\emptyset} > \mathsf{OBJ}$, then we *will not* perform this join. Instead, we *restart* the subproblem with a fresh inequality. In particular, set $\boldsymbol{\delta} = \boldsymbol{\delta} + w \cdot \mathbf{c}_{X,Y}$. Now we have $\delta_{Y|\emptyset} \geq w$. We restart the problem with the inequality

$$\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}', \mathbf{h} \rangle \tag{52}$$

satisfying the conditions stated in Lemma 4.10. We will show in the proof of Theorem 5.4 below that in this case $\|\boldsymbol{\lambda}\|_1 > w$, and the new inequality is a Shannon flow inequality with a smaller upper bound on the proof sequence length.

Lemma 5.3. *Let $X \subset Y \subseteq V$. Let $R(Y)$ be a table with $\log_2 |R| \leq n_{Y|\emptyset}$. Then, R can be partitioned into at most $k = 2 \log |R|$ sub-tables $R^{(1)}, \ldots, R^{(k)}$ such that $n^{(j)}_{Y|X} + n^{(j)}_{X|\emptyset} \leq n_{Y|\emptyset}$, for all $j \in [k]$, where*

$$\begin{aligned} n^{(j)}_{Y|X} &\stackrel{\text{def}}{=} \max_{\mathbf{t} \in \Pi_X(R^{(j)})} \log_2 \deg_{R^{(j)},Y}(\mathbf{t}) \\ n^{(j)}_{X|\emptyset} &\stackrel{\text{def}}{=} \log_2 |\Pi_X(R^{(j)})|. \end{aligned}$$

Proof. See Appendix C. □

Theorem 5.4. *Let $(\boldsymbol{\sigma}, \boldsymbol{\mu})$ be a witness of the guiding Shannon flow inequality, and D the minimum common denominator of all entries in $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$. PANDA runs in time $\tilde{O}(N + \text{poly}(OBJ) \cdot 2^{OBJ})$, where N is the input size.*

Proof. Consider an input proof sequence of length $\ell \leq D(3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1)$, thanks to Theorem 4.9. If we do not hit Case 4b, then there will be at most ℓ steps in the algorithm, where each step either takes $\tilde{O}(2^{\mathsf{OBJ}})$ time or spawns $O(\mathsf{OBJ})$ subproblems, for a total of $\tilde{O}(\text{poly}(\mathsf{OBJ}) \cdot 2^{\mathsf{OBJ}})$-time. A subproblem will terminate at producing a relation $T(\mathbf{A}_B)$ for some $B \in \mathcal{B}$ because the proof sequence will, by definition, reach a point where $\delta_{B|\emptyset} \geq w > 0$ for some $B \in \mathcal{B}$.

The worst case is obtained when the algorithm branches as far as possible only to have to restart at Case 4b with a slightly shorter proof sequence of length at most $D(3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1) - 1$. Thus, overall the exponent of OBJ (in $\text{poly}(\mathsf{OBJ})$) will be $\leq \frac{1}{2} D^2 (3\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\delta}\|_1 + \|\boldsymbol{\mu}\|_1)^2$. This constant is data-independent because the dual solution $(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ can be taken to be an extreme point of the dual polyhedron (44), whose constraints are only on the input query.

The last bit we have to verify is that every subproblem has a proof sequence that can eventually reach $\delta_{B|\emptyset} \geq w$ for some $B \in \mathcal{B}$. Since we start with a proof sequence and in Cases 1, 2, 3, and 4a we continue with the remaining steps of the input proof sequence, the only case we have to worry about proving the existence of a proof sequence reaching $\delta_{B|\emptyset} \geq w$ is Case 4b. In this case we start with inequality (52). Lemma 4.10

guarantees that $\lambda'_B > 0$ implies $\lambda_B > 0$. Hence, we will be done if we can show that $\|\boldsymbol{\lambda}'\|_1 > 0$. To this end, we show that the algorithm maintains another invariant, that for every subproblem it always holds that

$$\sum_{(X,Y)} n(\delta_{Y|X}) \le \|\boldsymbol{\lambda}\|_1 \cdot \mathsf{OBJ}. \tag{53}$$

We call the quantity $\sum_{(X,Y)} n(\delta_{Y|X})$ the "potential" of the subproblem at that point in time. At the initial iteration, the potential of the input problem was exactly the budget. When Case 1 applies, the potential is unchanged because $\delta_{I|I\cap J}$ was reduced by w, $\delta_{I\cup J|J}$ was increased by w, and they have the same support. When Case 2 applies, the potential does not increase because we subtracted $w \cdot n_{Y|\emptyset}$ from the potential and added $w \cdot n_{X|\emptyset} \le w \cdot n_{Y|\emptyset}$ instead. Similarly in Case 3, we subtracted $w \cdot n_{Y|\emptyset}$ from the potential, and added $w \cdot (n^{(j)}_{X|\emptyset} + n^{(j)}_{Y|X}) \le w \cdot n_{Y|\emptyset}$ to the potential. In Case 4a, we subtracted $w \cdot (n_{X|\emptyset} + n_{W|Z})$ and added at most the same amount to the potential. In Case 4b, since $\boldsymbol{\delta}' \le \boldsymbol{\delta}$ and $\delta'_{Y|\emptyset} = \delta_{Y|\emptyset} - w$, the potential is reduced by strictly more than $w \cdot \mathsf{OBJ}$ because we subtracted $w \cdot (n_{X|\emptyset} + n_{W|Z}) > w \cdot \mathsf{OBJ}$ from the potential. Since $\|\boldsymbol{\lambda}\|_1$ was reduced by at most w, invariant (53) still holds. This also proves that $\|\boldsymbol{\lambda}\|_1 > w$, because the residual potential, which is non-negative, is strictly smaller than $(\|\boldsymbol{\lambda}\|_1 - w)\mathsf{OBJ} \le \|\boldsymbol{\lambda}'\|_1 \mathsf{OBJ}$.

Finally, we get back to assumption (50) and how to enforce it in case it is not initially satisfied. Suppose there is an input relation $T(\mathbf{A}_Y)$ where $n_{Y|\emptyset} \stackrel{\text{def}}{=} \log_2 |T| > \mathsf{OBJ}$. If $\delta_{Y|\emptyset} = 0$, then T does not contribute to the bound/proof sequence and can be ignored. If $\delta_{Y|\emptyset} > 0$, then we could have replaced the original inequality with a new one

$$\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$$

satisfying the conditions stated in Lemma 4.10, in the same way we handled Case 4b above. □

6 Width parameters for queries with degree constraints

This section shows that known width parameters can be unified under the same simple umbrella, making it easy to generalize them to the case of queries with degree constraints.

6.1 A slight reinterpretation of existing width parameters

We first reformulate existing width parameters under a common framework. Recall from Section 2.1 that there are two classes of width parameters: the first class captures the class of algorithms which seek the best tree decomposition with the worst bag runtime, while the second class captures the class of algorithms which adapts the tree decomposition to the instance at hand.

Definition 6.1 (Widths under degree constraints)**.** Let $\mathcal{F}$ denote a class of non-negative set functions on V. The *$\mathcal{F}$-minimax width* of a query Q with hypergraph $\mathcal{H}$ is defined as

$$\mathsf{Minimaxwidth}_{\mathcal{F}}(Q) \stackrel{\text{def}}{=} \min_{(T,\chi)} \max_{t \in V(T)} \sup_{h \in \mathcal{F}} h(\chi(t)). \tag{54}$$

The *$\mathcal{F}$-maximin width* of $\mathcal{H}$ is defined as

$$\mathsf{Maximinwidth}_{\mathcal{F}}(Q) \stackrel{\text{def}}{=} \sup_{h \in \mathcal{F}} \min_{(T,\chi)} \max_{t \in V(T)} h(\chi(t)). \tag{55}$$

(Since we assume nothing about the topology of $\mathcal{F}$, we have to use sup instead of max.)

We have seen several classes of set functions capturing relaxation or specializations of the entropic functions $(\mathsf{M}_n, \overline{\Gamma}^*_n, \Gamma_n, \mathsf{SA}_n)$. We have also seen other classes of set functions capturing constraints one can infer from the statistics of input relations $(\mathsf{VD}, \mathsf{ED}, \mathsf{HDC}\ldots)$.

The following simple propositions explain why the minimax and maximin widths capture all width parameters we discussed in Section 2. A proof can be found in the appendix. Note the following obvious inclusion relationship:

$$\mathsf{M}_n \subset \overline{\Gamma}^*_n \subset \Gamma_n \subset \mathsf{SA}_n.$$

When we maximize a set function over one of these classes to obtain a bound, we have a hierarchy of bounds. A proof of the following simple lemma can be found in the appendix.

Lemma 6.2. *If $\mathcal{G} \subseteq \mathcal{F}$ are two classes of functions, then* $\mathsf{Minimaxwidth}_{\mathcal{G}}(\mathcal{H}) \leq \mathsf{Minimaxwidth}_{\mathcal{F}}(\mathcal{H})$ *and* $\mathsf{Maximinwidth}_{\mathcal{G}}(\mathcal{H}) \leq \mathsf{Maximinwidth}_{\mathcal{F}}(\mathcal{H})$. *For a fixed class $\mathcal{F}$ of functions, we have* $\mathsf{Maximinwidth}_{\mathcal{F}}(\mathcal{H}) \leq \mathsf{Minimaxwidth}_{\mathcal{F}}(\mathcal{H})$.

In particular, a bound on SA_n is probably too weak, and a bound using M_n is probably too strong. However, depending on the constraints we want to impose, some parts of the hierarchy collapse.

Proposition 6.3. *Let Q be a conjunctive query with no FDs whose hypergraph is $\mathcal{H} = (V, \mathcal{E})$. Then the followings hold (Recall notation from Section 2):*

$$1 + \mathsf{tw}(\mathcal{H}) = \mathsf{Minimaxwidth}_{\mathcal{F} \cap \mathsf{VD}(\mathbf{1}_V)}(Q) = \mathsf{Maximinwidth}_{\mathcal{F} \cap \mathsf{VD}(\mathbf{1}_V)}(Q), \quad \forall \mathcal{F} \in \{\mathsf{M}_n, \overline{\Gamma}^*_n, \Gamma_n, \mathsf{SA}_n\} \tag{56}$$

$$\mathsf{ghtw}(\mathcal{H}) = \mathsf{Minimaxwidth}_{\mathsf{SA}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})}(Q) = \mathsf{Maximinwidth}_{\mathsf{SA}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})}(Q) \tag{57}$$

$$\mathsf{fhtw}(\mathcal{H}) = \mathsf{Minimaxwidth}_{\mathcal{F} \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})}(Q), \quad \forall \mathcal{F} \in \{\mathsf{M}_n, \overline{\Gamma}^*_n, \Gamma_n\} \tag{58}$$

$$\mathsf{subw}(\mathcal{H}) = \mathsf{Maximinwidth}_{\Gamma_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})}(Q) \tag{59}$$

$$\mathsf{adw}(\mathcal{H}) = \mathsf{Maximinwidth}_{\mathsf{M}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})}(Q). \tag{60}$$

Note that the use of all-1 vectors is to normalize the widths so that they are data-size independent. To make them more data-size adaptive, we should use the corresponding input relation sizes, in the same way the AGM-bound is more adaptive than the fractional edge cover number bound.

From the above lemmas and the chains of inclusion $M_n \subset \overline{\Gamma}^*_n \subset \Gamma_n \subset \mathsf{SA}_n$ and $\mathsf{ED}(\mathbf{1}_\mathcal{E}) \subseteq \mathsf{VD}(\mathbf{1}_V)$, the following follows trivially

Corollary 6.4. *Let Q be a conjunctive query (with no FDs nor degree constraints) whose hypergraph is $\mathcal{H}$, then*

$$1 + \mathsf{tw}(\mathcal{H}) \geq \mathsf{ghtw}(\mathcal{H}) \geq \mathsf{fhtw}(\mathcal{H}) \geq \mathsf{subw}(\mathcal{H}) \geq \mathsf{adw}(\mathcal{H}) \tag{61}$$

And, the gap between any two consecutive entries in the above series is unbounded; namely, for any two consecutive entries and any constant there is a query for which the difference between the entries is greater than the given constant.

Note from (57) the interesting fact that the set SA_n is so large that switching to the Maximinwidth does not help reduce the objective.

Example 6.5 (Unbounded gap between fhtw and subw)**.** To further illustrate the strength of the minimax/maximin characterization, we use the above characterization to present an example where the gap between fhtw and subw is unbounded. (Marx [34,35] constructed a hypergraph where subw is bounded and fhtw is unbounded; however, his example is much more involved.) Consider a query whose graph $\mathcal{H} = (V, \mathcal{E})$ and is defined as follows. The vertex set $V = I_1 \cup I_2 \cup \cdots \cup I_{2k}$ is a disjoint union of $2k$ sets of vertices. Each set I_j has m vertices in it. There is no edge between any two vertices within the set I_j for every $j \in [2k]$, namely I_j is an independent set. The edge set $\mathcal{E}$ of the hypergraph is the union of $2k$ complete bipartite graphs $K_{m,m}$:

$$\mathcal{E} \stackrel{\text{def}}{=} I_1 \times I_2 \cup I_2 \times I_3 \cup \cdots \cup I_{2k-1} \times I_{2k} \cup I_{2k} \times I_1.$$

To each edge in $\mathcal{E}$ there corresponds an input relation. We first bound the fhtw of this graph. The neighborhood of every vertex contains $2m + 1$ vertices: the vertex and its neighboring two independent sets for a total of $2m$ independent vertices. Since we need one edge to cover every vertex of an independent set, the best fractional cover bound for this neighborhood is $2m$, namely $\mathsf{fhtw}(\mathcal{H}) \geq 2m$.

Next, we bound $\mathcal{H}$'s submodular width. Let h be any submodular function.

- Case 1: $h(I_i) \le \theta$ for some $i \in [2k]$. WLOG assume $h(I_1) \le \theta$. Consider the tree decomposition

  For bag $B = I_1 \cup I_i \cup I_{i+1}$,

$$h(B) \le h(I_1) + h(I_i \cup I_{i+1}) \le \theta + m.$$

- Case 2: $h(I_i) > \theta$ for all $i \in [2k]$. Consider the tree decomposition

  From submodularity, it is easy to see that

$$\begin{aligned} h(B_1) &\le km - (k-1)\theta \\ h(B_2) &\le km - (k-1)\theta \end{aligned}$$

  Thus, by setting $\theta = (1 - 1/k)m$ we have just proved that $\mathsf{subw}(\mathcal{H}) \le m(2 - 1/k)$. By increasing m, the gap between $m(2 - 1/k)$ and $2m$ is infinite for a fixed $k \ge 2$.

6.2 New width parameters

To define a width parameter, one takes the minimax or maximin width of a function class $\mathcal{F}$. The function class is the intersection of two function classes: (1) an approximation of $\overline{\Gamma}^*_n$, and (2) the set of constraints we obtain from the input database instance. This section extends the traditional width parameters in two ways. First, as shown in Section 3.1 we know that there is a gap between the polymatroid bound and the entropic bound; and hence it is natural to use $\overline{\Gamma}^*_n$ itself instead of some approximation of it. Second, as aforementioned the degree constraints strictly generalize both cardinality constraints and FD constraints.

Definition 6.6. We define the following width parameters for queries Q with degree constraints DC. The first two parameters are generalizations of fhtw and subw under degree constraints, and the last two are their entropic versions:

$$\mathsf{da\text{-}fhtw}(Q) \stackrel{\text{def}}{=} \mathsf{Minimaxwidth}_{\Gamma_n \cap \mathsf{HDC}}(Q) \tag{62}$$

$$\mathsf{da\text{-}subw}(Q) \stackrel{\text{def}}{=} \mathsf{Maximinwidth}_{\Gamma_n \cap \mathsf{HDC}}(Q) \tag{63}$$

$$\mathsf{eda\text{-}fhtw}(Q) \stackrel{\text{def}}{=} \mathsf{Minimaxwidth}_{\overline{\Gamma}^*_n \cap \mathsf{HDC}}(Q) \tag{64}$$

$$\mathsf{eda\text{-}subw}(Q) \stackrel{\text{def}}{=} \mathsf{Maximinwidth}_{\overline{\Gamma}^*_n \cap \mathsf{HDC}}(Q). \tag{65}$$

Note that the bounds in the constraints HDC are not normalized as in the sets $\mathsf{ED}(\mathbf{1}_{\mathcal{E}})$ or $\mathsf{VD}(\mathbf{1}_V)$ in the traditional width parameters. This is because normalizing makes it less general than it can be, and it does not make practical sense to assume that all degree bounds are the same! (For example, the FD-based degree bounds are always 0, while the relation-size-based degree bounds are $\log_2 N_F$.) Consequently, we will use the $\log_2 N$ scaled up versions of the traditional width parameters to compare with our new width parameters.

The following relationships hold between these four quantities.

Proposition 6.7. *For any query Q with degree constraints*

$$\begin{array}{ccc} \mathsf{eda\text{-}subw}(Q) & \le & \mathsf{eda\text{-}fhtw}(Q) \\ \mathsf{I}\wedge & & \mathsf{I}\wedge \\ \mathsf{da\text{-}subw}(Q) & \le & \mathsf{da\text{-}fhtw}(Q). \end{array}$$

The quantities $\mathsf{eda\text{-}fhtw}(Q)$ *and* $\mathsf{da\text{-}subw}(Q)$ *are not comparable. The gap between the two sides of any of the above four inequalities can be made arbitrarily large by some query* Q *and corresponding input database instance.*

A proof of the proposition can be found in Appendix D. In this paper, we show that our algorithm PANDA can achieve a runtime of $\tilde{O}(2^{\mathsf{da\text{-}fhtw}(Q)} \cdot \mathrm{poly}(\log N) + |\mathsf{output}|)$ *and* $\tilde{O}(2^{\mathsf{da\text{-}subw}(Q)} \cdot \mathrm{poly}(\log N) + |\mathsf{output}|)$. Another somewhat interesting observation which follows from the above is the following. Due to the fact that every non-negative modular set function is entropic, we have

Corollary 6.8. *When* Q *has only cardinality bounds normalized to be* $\mathsf{ED}(\mathbf{1}_{\mathcal{E}})$ *(i.e. no FD nor proper degree bounds), we have* $\mathsf{adw}(Q) \leq \mathsf{eda\text{-}subw}(Q)$.

In particular, from Marx [34, 35] it follows that for these queries we have $\mathsf{da\text{-}subw}(Q) = \mathsf{subw}(Q) = O(\mathsf{adw}^4(Q)) = O(\mathsf{eda\text{-}subw}^4(Q))$. Thus, when there is no FD and proper degree bounds, there is no query Q for which eda-subw is bounded and da-subw is unbounded. It is open whether or not the same relationship holds when Q has FDs and/or degree bounds.

Example 6.9. Consider the 4-cycle query C_4 shown in equation (10), which has no proper degree bounds (only input size bounds). Let N be the common upperbound on all input relations' sizes, we will show that $\mathsf{da\text{-}subw}(C_4) = \mathsf{eda\text{-}subw}(C_4) = 3/2 \log N$ and $\mathsf{da\text{-}fhtw} = \mathsf{eda\text{-}fhtw}(C_4) = 2 \log N$. Since there is no degree bound, $\mathsf{da\text{-}fhtw}(C_4) = \mathsf{fhtw}(C_4) \cdot \log N = 2 \log N$, and $\mathsf{da\text{-}subw}(C_4) = \mathsf{subw}(C_4) \cdot \log N \leq 3/2 \log N$ (see Example 6.5). To show $\mathsf{da\text{-}subw}(C_4) \geq 3/2 \log N$, consider the function $\bar{h}(F) = \frac{|F|}{2} \log N$, for all $F \subseteq [4]$. This function is in Γ_4^*; and, for any bag B of the two tree decompositions shown in Figure 2, $h(B) = 3/2 \log N$. Hence, $3/2 \log N \geq \mathsf{da\text{-}subw}(C_4) \geq \mathsf{eda\text{-}subw}(C_4) \geq 3/2 \log N$. The fact that $\mathsf{eda\text{-}fhtw}(C_4) = 2 \log N$ can be shown similarly.

We have mentioned quite a few known and proved new bounds in this paper. The bounds can be summarized systematically as follows. Each bound is identified by coordinates (X, Y, Z).

- The X-axis represents the entropy approximation that is being used: one starts from the desired target $\overline{\Gamma}_n^*$, then relax it to Γ_n and SA_n. The inclusion chain is $\overline{\Gamma}_n^* \subset \Gamma_n \subset \mathsf{SA}_n$.
- The Y-axis represents the constraints we can extract from the input database instance, where we can go from bounding domain sizes, relation sizes, to incorporating more refined degree bounds and functional dependencies. One chain of inclusion is

$$\mathsf{VD}(\log N \cdot \mathbf{1}_V) \supset \mathsf{ED}(\log N \cdot \mathbf{1}_V) \supset \mathsf{ED}((\log N_F)_{F \in \mathcal{E}}) \supset \mathsf{HDC}.$$

- The Z-axis represents the level of sophistication of the query plan that is being considered in this bound. The simplest query plan just joins everything together without computing any tree decomposition – or, equivalently, this is the plan that uses the trivial tree decomposition with one bag. (We still need PANDA even for this query plan, because no known algorithm can achieve the degree-aware output size bounds.) Then, one can get more sophisticated with computing a tree decomposition before computing the bags. And, lastly the query plan can also be adaptive to the input instance, yielding the submodular-width style of complexity.

The bounds are summarized in Figure 4. If bound A has coordinates (X_1, Y_1, Z_1) that are smaller than the coordinates (X_2, Y_2, Z_2) of bound B (i.e. $X_1 \leq X_2 \wedge Y_1 \leq Y_2 \wedge Z_1 \leq Z_2$), then bound $A \leq$ bound B.

6.3 Achieving degree-aware width parameters with PANDA

A basic fact that we employ in this section is the following: given an α-acylic query (with no FD nor degree constraints), the query can be computed in time $\tilde{O}(|\mathsf{input}| + |\mathsf{output}|)$ [22, 44]. In particular, let (T, χ) be some tree decomposition, and suppose we have a query for which every bag B of the tree decomposition corresponds uniquely to an input relation R_B; then, the query is α-acyclic (the tree decomposition is the

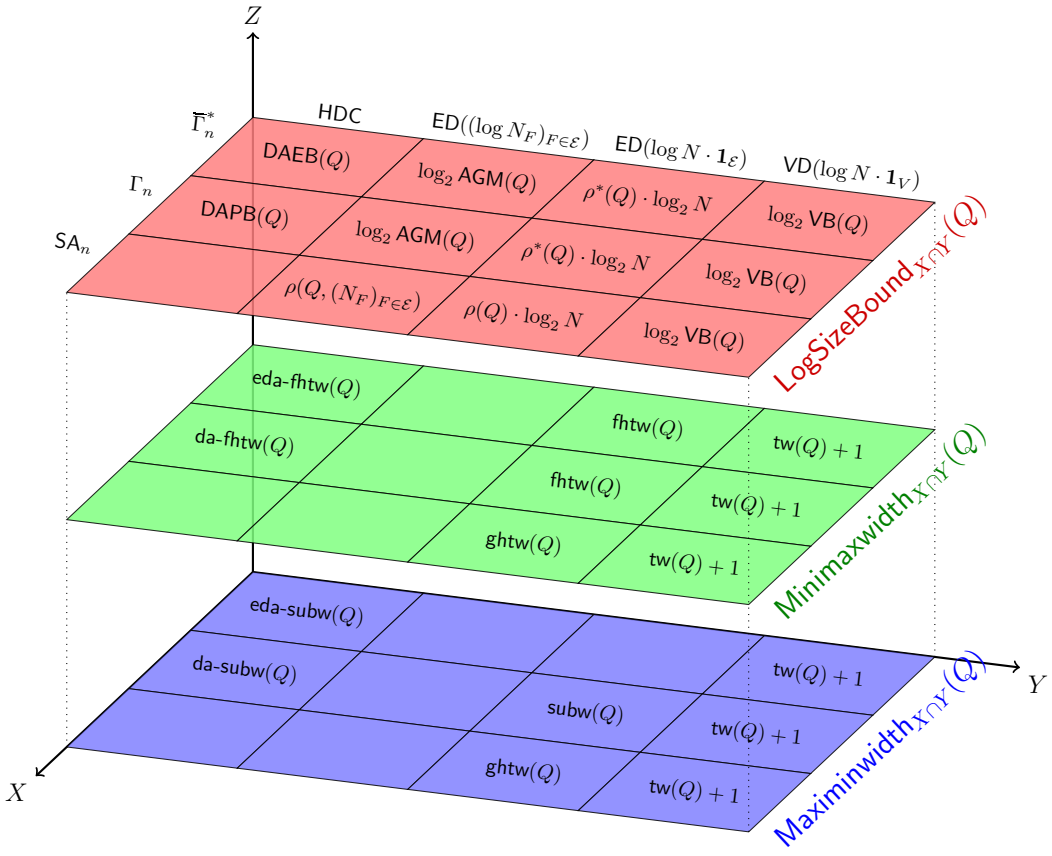

Figure 4: The hierarchy of bounds we discussed in this paper

join tree of the query), and it can be computed in linear time in the input plus output sizes (modulo a log factor).

With increasing levels of complexity, we explain how PANDA can be used to evaluate a (full or Boolean) conjunctive query achieving (1) the degree-aware polymatroid size bound defined in (26), (2) the degree-aware version of the fractional hypertree width defined in (62), and (3) the degree-aware version of the submodular width defined in (63).

6.3.1 Conjunctive query with degree constraints

Let $B \subseteq V$ be any fixed set. Let $\mathbf{e}^B = (e(Y|X))_{(X,Y)\in\mathcal{P}} \in \mathbb{Q}_+^{\mathcal{P}}$ denote the unit vector where $e(B|\emptyset) = 1$ and $e(Y|X) = 0$ otherwise. Then, inequality (68) is $\langle \mathbf{e}^V, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$, and inequality (71) is $\langle \mathbf{e}^B, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}, \mathbf{h}\rangle$.

Consider a full conjunctive query Q with degree constraints HDC. Since a full conjunctive query is just the disjunctive datalog query (5) with $\mathcal{B} = \{V\}$, we have an immediate corollary. However, we spell out more details than necessary here for the reader to get a sense of how the Shannon flow inequality looks in this simple case. The bound $\mathsf{DAPB}(Q)$ is represented by the LP in (26), which can be written more concretely as

$$\begin{array}{lrll} \max & h(V) & & \\ \text{s.t.} & h(Y) - h(X) \le n_{Y|X} & (X,Y,n_{Y|X}) \in \mathsf{DC} & \text{(degree constraints)} \\ & h(I\cup J|J) - h(I|I\cap J) \le 0, & I \perp J & \text{(submodularity)} \\ & h(Y) - h(X) \ge 0, & \emptyset \neq X \subset Y \subseteq V & \text{(monotonicity)} \\ & h(Z) \ge 0, & \emptyset \neq Z \subseteq V. & \text{(non-negativity)} \end{array} \tag{66}$$

The dual of (66) is

$$\begin{array}{lll} \min & \sum_{(X,Y)\in\mathsf{DC}} n_{Y|X}\cdot\delta_{Y|X} & \\ \text{s.t.} & \mathsf{inflow}(V) & \ge 1 \\ & \mathsf{inflow}(Z) & \ge 0, \emptyset \neq Z \subset V. \\ & (\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu}) & \ge \mathbf{0}. \end{array} \tag{67}$$

Let h^* denote an optimal solution to (66), and $(\boldsymbol{\delta}^*, \boldsymbol{\sigma}^*, \boldsymbol{\mu}^*)$ a dual-optimal solution, then $h^*(V) = \sum_{(X,Y)\in\mathsf{DC}} \delta^*_{Y|X}\cdot n_{Y|X}$. The $\mathsf{DAPB}(Q)$ bound can also be written as $2^{h^*(V)} = \prod_{(X,Y)\in\mathsf{DC}} 2^{n_{Y|X}\delta^*_{Y|X}}$. This was also the form of $\mathsf{AGM}(Q)$ for queries with no degree constraints. (See also Proposition 6.3 and its proof in Appendix D). In this case, Proposition 4.3 states the following. Given $\boldsymbol{\delta} \in \mathbb{Q}_+^{\mathsf{DC}}$, the inequality

$$h(V) \le \sum_{(X,Y)\in\mathsf{DC}} \delta_{Y|X}\cdot h(Y|X) \tag{68}$$

is a Shannon flow inequality if and only if there exist $\boldsymbol{\sigma}$ and $\boldsymbol{\mu}$ such that $(\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu})$ is feasible to the dual LP (67). (In particular, the inequality holds when $\boldsymbol{\delta} = \boldsymbol{\delta}^*$.) Note that inequality (68) implies the upperbound on $h(V)$ we wanted because $h(Y|X) \le n_{Y|X}$. Note also that inequality (68) has only one target V.

Corollary 6.10. *A query Q with degree constraints DC can be solved by PANDA in time*

$$\tilde{O}(N + \mathrm{poly}(\log N)\cdot 2^{\mathsf{LogSizeBound}_{\Gamma_n\cap\mathsf{HDC}}(Q)}).$$

Proof. Recall that $\mathsf{LogSizeBound}_{\Gamma_n\cap\mathsf{HDC}}(Q) = \mathsf{OPT}$, where OPT is the optimal objective value of the LP (66), whose dual is (67). Let $(\boldsymbol{\delta}^*, \boldsymbol{\sigma}^*, \boldsymbol{\mu}^*)$ denote a dual-optimal solution. Then, from Proposition 4.3 $\langle \mathbf{e}^V, \mathbf{h}\rangle \le \langle \boldsymbol{\delta}^*, \mathbf{h}\rangle$ is a Shannon flow inequality. When feeding this Shannon flow inequality to PANDA, the quantity OBJ defined in (49) is exactly OPT due to strong duality. Moreover, since $\mathsf{OPT} \le \log_2 \mathsf{VB}(Q)$, we have $\mathsf{OBJ} \le |V|\cdot\log_2 N = O(\log_2 N)$. The output of PANDA is basically a single table T_V which *contains* Q. To get Q exactly, we semijoin-reduce T_V with every input relation R_F, $F \in \mathcal{E}$; namely we set $T_V = T_V \ltimes R(F)$ for all $F \in \mathcal{E}$. ∎

6.3.2 Achieving the degree-aware fractional hypertree width

Next let us consider the degree-aware version of the fractional hypertree width defined in (62), whose full definition is

$$\mathsf{da\text{-}fhtw}(Q) \stackrel{\text{def}}{=} \min_{(T,\chi)} \max_{t\in V(T)} \max_{h\in\Gamma_n\cap\mathsf{HDC}} h(\chi(t)).$$

Suppose we want to compute da-fhtw, a bruteforce algorithm is to go over all tree decompositions (T,χ) of $\mathcal{H}$[14]; for each (T,χ) we take a bag $B=\chi(t)$ and solve the inner most optimization problem [15]

$$\max\{h(B) \mid h\in\Gamma_n\cap\mathsf{HDC}\} \tag{69}$$

(Note a subtle fact that in the case of no degree constraints, this LP can be reduced to the LP (66) by restricting all functions down to the universe B; but this does not hold when we add FDs to the query, and it is certainly not true when we add arbitrary degree constraints.) The dual of (69) is the following LP:

$$\begin{array}{lll} \min & \sum_{(X,Y)\in\mathsf{DC}} n_{Y|X}\cdot\delta_{Y|X} & \\ \text{s.t.} & \mathsf{inflow}(B) & \geq 1 \\ & \mathsf{inflow}(Z) & \geq 0, \emptyset\neq Z\subseteq V. \\ & (\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu}) & \geq \mathbf{0}. \end{array} \tag{70}$$

And, the corresponding Shannon flow inequality is the following, with one target B: given $\boldsymbol{\delta}\in\mathbb{Q}_+^{\mathsf{DC}}$, the inequality

$$h(B) \leq \sum_{(X,Y)\in\mathsf{DC}} \delta_{Y|X}\cdot h(Y|X) \tag{71}$$

is a Shannon flow inequality if and only if there exist $\boldsymbol{\sigma}$ and $\boldsymbol{\mu}$ such that the vector $(\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu})$ is feasible to the dual LP (70).

Corollary 6.11. *A query Q with degree constraints DC can be solved by PANDA in time*

$$\tilde{O}(N+\mathrm{poly}(\log N)\cdot 2^{\mathsf{da\text{-}fhtw}(Q)}+|\mathsf{output}|).$$

Proof. As discussed above, a tree decomposition (T,χ) of Q for which $\max_{t\in V(T)}\max_{h\in\Gamma_n\cap\mathsf{HDC}} h(\chi(t)) = \mathsf{da\text{-}fhtw}(Q)$ can be computed bruteforcely by going through at most $n!$ tree decompositions and solving a total of $\leq 2^n$ distinct linear programs, each of which is data independent (except for a log-factor in the data to compute the degree bounds in advance).

Now suppose we have already fixed an optimal tree decomposition (T,χ). For every bag B of the tree decomposition, we use PANDA to compute a relation T_B for which $\Pi_B(Q)\subseteq T_B$. After that we set $T_B = T_B\ltimes R_F$ for every $F\in\mathcal{E}$. (By definition of a tree decomposition, for every input relation R_F, there must be a bag $B\supseteq F$.) Finally, Q is the join of all these tables T_B, which is now an α-acyclic query solvable in linear time (input plus output size) by Yannakakis's algorithm.

For a fixed bag B of the optimal tree decomposition, let OPT be the optimal objective value of the LP (69). Let $(\boldsymbol{\delta}^*,\boldsymbol{\sigma}^*,\boldsymbol{\mu}^*)$ denote a dual-optimal solution. Then, from Proposition 4.3 we know $\langle\mathbf{e}^B,\mathbf{h}\rangle\leq\langle\boldsymbol{\delta}^*,\mathbf{h}\rangle$ is a Shannon flow inequality. When feeding this Shannon flow inequality to PANDA, the quantity OBJ defined in (49) is exactly OPT due to strong duality. ∎

[14] There are at most $n!$ non-redundant tree decompositions, each of which has at most n bags (Proposition 2.3).

[15] The total number of distinct optimization problems that we have to solve for all tree decompositions is $\leq 2^n$, since we have one problem for each distinct choice of B.

6.3.3 Achieving the degree-aware submodular width

We next investigate the degree-aware submodular width defined in (63). In order to compute this width, even in a bruteforce manner, we need an auxiliary lemma, which is somewhat related to Neumann's minimax theorem [17].

Lemma 6.12. *Let A and B be two finite sets, and $f : A \times B \to \mathbb{R}$ be any function. Let B^A denote the set of all maps from A to B. Then, the following holds:*

$$\min_{a\in A}\max_{b\in B} f(a,b) = \max_{\beta\in B^A}\min_{a\in A} f(a,\beta(a)).$$

Intuitively, on the LHS we select for each $a \in A$ a neighbor b for which $f(a,b)$ is maximized; call such neighbor a's "representative". Then, we select the a with the least-weight representative. On the RHS, we have a "representative selector" β; we pick the a-value with the least-weight selected representative, and then maximize over all selectors.

Corollary 6.13. *A query Q with degree constraints DC can be solved by PANDA in time*

$$\tilde{O}(N + \mathrm{poly}(\log N)\cdot 2^{\mathsf{da\text{-}subw}(Q)} + |\mathsf{output}|).$$

Proof. We first apply Lemma 6.12 to reformulate (63). To this end, we need a few notations. Let $\mathcal{M}$ be the set of all maps $\beta : \mathsf{TD} \to 2^V$, such that $\beta(T,\chi) = \chi(t)$ for some $t \in V(T)$. In English, β is a "bag selector" map that picks out a bag from each tree decomposition (T,χ). Let $\mathbf{B}$ be the collection of images of all $\beta \in \mathcal{M}$, i.e.

$$\mathbf{B} = \{\mathcal{B} \mid \mathcal{B} = \mathsf{image}(\beta) \text{ for some } \beta \in \mathcal{M}\}. \tag{72}$$

Using Lemma 6.12, we can rewrite (63) as follows.

$$\begin{aligned}
\mathsf{da\text{-}subw}(\mathcal{H}) &= \max_{h\in\Gamma_n\cap\mathsf{HDC}}\min_{(T,\chi)\in\mathsf{TD}}\max_{t\in V(T)} h(\chi(t))\\
\text{(Lemma 6.12)} &= \max_{h\in\Gamma_n\cap\mathsf{HDC}}\max_{\beta\in\mathcal{M}}\min_{(T,\chi)} h(\beta(T,\chi))\\
&= \max_{h\in\Gamma_n\cap\mathsf{HDC}}\max_{\beta\in\mathcal{M}}\min_{B\in\mathsf{image}(\beta)} h(B)\\
&= \max_{\beta\in\mathcal{M}}\max_{h\in\Gamma_n\cap\mathsf{HDC}}\min_{B\in\mathsf{image}(\beta)} h(B)\\
&= \max_{\mathcal{B}\in\mathbf{B}}\max_{h\in\Gamma_n\cap\mathsf{HDC}}\min_{B\in\mathcal{B}} h(B)\\
\text{(Lemma 4.2)} &= \max_{\mathcal{B}\in\mathbf{B}}\underbrace{\max_{h\in\Gamma_n\cap\mathsf{HDC}}\left\{\sum_{B\in\mathcal{B}}\lambda_B h(B)\right\}}_{\text{Linear program (41)}}
\end{aligned} \tag{73}$$

In (73), for a fixed $\mathcal{B} \in \mathbf{B}$ the inner max is exactly LP (41) whose dual is (42). In particular, to compute the $\mathsf{da\text{-}subw}(Q)$, we can solve a collection of linear programs and take the maximum solution among them. Since there is a different linear program for each valid choice of $\mathcal{B}$, the total number of linear programs is $\le 2^{2^{|V|}}$.

In order to compute Q in the desired time, we mimic this strategy in the algorithm. For each $\mathcal{B} \in \mathbf{B}$, we solve the LP (41). Let $(\boldsymbol{\delta}^*, \boldsymbol{\sigma}^*, \boldsymbol{\mu}^*)$ denote a dual optimal solution. From Proposition 4.3 $\langle\boldsymbol{\lambda}, \mathbf{h}\rangle \le \langle\boldsymbol{\delta}^*, \mathbf{h}\rangle$ is a Shannon flow inequality. On this input PANDA computes a tuple $\mathbf{T}_{\mathcal{B}} = (T_B)_{B\in\mathcal{B}}$ of tables such that, for every $\mathbf{a} \in Q$ there exists a $B \in \mathcal{B}$ for which $\Pi_B(\mathbf{a}) \in T_B$.

Let $M = |\mathbf{B}|$ and suppose $\mathbf{B} = \{\mathcal{B}_1, \ldots, \mathcal{B}_M\}$. We prove the following claims:

Claim 1: for every $(B_1, \ldots, B_M) \in \prod_{i=1}^M \mathcal{B}_i$, there is a tree decomposition $(T, \chi) \in \mathsf{TD}(Q)$ such that, for every tree node $t \in V(T)$, $\chi(t) = B_j$ for some $j \in [M]$. Breaking ties arbitrarily, we call this tree decomposition *the* tree decomposition (of Q) associated with the tuple $(B_1, \ldots, B_M)$.

Claim 2: for any tuple $(B_1, \ldots, B_M) \in \prod_{i=1}^M \mathcal{B}_i$ with associated tree decomposition (T, χ), define $J(B_1, \ldots, B_M) \stackrel{\text{def}}{=} \bowtie_{t \in V(T)} T_{\chi(t)}$. Then,

$$Q \subseteq \left(\bigcup_{(B_1,\ldots,B_M) \in \prod_{i=1}^M \mathcal{B}_i} \bowtie_{j=1}^M T_{B_j} \right) \subseteq \left(\bigcup_{(B_1,\ldots,B_M) \in \prod_{i=1}^M \mathcal{B}_i} J(B_1, \ldots, B_M) \right). \tag{74}$$

Assuming the claims, the query can be computed by running Yannakakis algorithm to compute all the relations $J(B_1, \ldots, B_M)$ within a runtime of $\tilde{O}(2^{\mathsf{da\text{-}subw}(Q)} + |J(B_1, \ldots, B_M) \cap Q|)$. If we apply Yannakakis's algorithm straightup on $J(B_1, \ldots, B_M)$, then we can attain the runtime $\tilde{O}(2^{\mathsf{da\text{-}subw}(Q)} + |J(B_1, \ldots, B_M)|)$, because every table T_{B_j} has size bounded by $2^{\mathsf{da\text{-}subw}(Q)}$. To reduce the runtime to $\tilde{O}(2^{\mathsf{da\text{-}subw}(Q)} + |J(B_1, \ldots, B_M) \cap Q|)$, we semijoin-reduce every table T_{B_j} with every input relation. There are $\prod_{i=1}^M |\mathcal{B}_i|$ such problems, which is a query-complexity quantity.

We next prove Claim 1. Fix a tuple $(B_1, \ldots, B_M) \in \prod_{i=1}^M \mathcal{B}_i$. Suppose to the contrary that for *every* tree decomposition (T, χ) there is a tree node $t \in V(T)$ such that $\chi(t) \neq B_j$ for every $j \in [M]$. Call the bag $\chi(t)$ a *missed* bag of the tree decomposition. Consider a bag selector $\beta : \mathsf{TD}(Q) \to 2^V$ where $\beta(T, \chi)$ is exactly the missed bag of the tree decomposition (T, χ). Note that by definition of $\mathbf{B}$ we have $\mathsf{image}(\beta) = \mathcal{B}_k$ for some $k \in [M]$. This is a contradiction because B_k must then be the missed bag of some tree decomposition, but it is not missed anymore.

Finally, we prove Claim 2. Consider an output tuple $\mathbf{a} \in Q$. For each $j \in [M]$, let B_j denote the bag for which $\Pi_{B_j}(\mathbf{a}) \in T_{B_j}$. Then, obviously $\mathbf{a} \in J(B_1, \ldots, B_M)$. This proves the first inclusion in (74). The second inclusion is obvious because the join $J(B_1, \ldots, B_M)$ drops some tables from the join $\bowtie_{j=1}^M T_{B_j}$. □

Example 6.14. Consider the 4-cycle query C_4 shown in (10) with two tree decompositions shown in Figure 2. Suppose the bag selector function chose $\beta(\mathrm{TD}_1) = \{1, 2, 3\}$ and $\beta(\mathrm{TD}_2) = \{2, 3, 4\}$. So $\mathcal{B} = \{\{1, 2, 3\}, \{2, 3, 4\}\}$. Our linear program is

$$\max\big\{\min\left(h(\{1, 2, 3\}), h(\{2, 3, 4\})\right) \mid h \in \Gamma_4 \cap \mathsf{HDC}\big\}.$$

This is connected to a Shannon flow inequality of the form (43) as follows.

$$\begin{aligned}
\min\{h(\{1, 2, 3\}), h(\{2, 3, 4\})\} &\leq \frac{1}{2}h(\{1, 2, 3\}) + \frac{1}{2}h(\{2, 3, 4\}) \\
&\leq \frac{1}{2}(h(\{1, 2\}) + h(\{2, 3\}) - h(\{2\})) + \frac{1}{2}(h(\{2, 3\}) + h(\{3, 4\}) - h(\{3\})) \\
&= \frac{1}{2}(h(\{1, 2\}) + h(\{2, 3\}) + h(\{3, 4\})) + \frac{1}{2}(h(\{2, 3\}) - h(\{2\}) - h(\{3\})) \\
&\leq \frac{1}{2}(h(\{1, 2\}) + h(\{2, 3\}) + h(\{3, 4\}).
\end{aligned}$$

The corresponding Shannon flow inequality is

$$\frac{1}{2}h(\{1, 2, 3\}) + \frac{1}{2}h(\{2, 3, 4\}) \leq \frac{1}{2}(h(\{1, 2\}) + h(\{2, 3\}) + h(\{3, 4\})).$$

6.4 Using PANDA to answer aggregate queries

Now instead of a conjunctive query, suppose we were dealing with a query with aggregates. In particular, let the input be an FAQ-SS query [31] as described earliear in Section 2.4. We can still follow the same outline of Sections 6.3.2 and 6.3.3 to answer such a query within basically the same width bounds. The main difference is that instead of computing a relation T_B for every bag B of the tree decomposition at hand, we

will be computing a factor ψ_B. And instead of applying Yannakakis's algorithm on those factors, we can apply InsideOut algorithm from [31].

One subtle issue here is the decomposition step of PANDA described earlier in Section 5. In particular, when PANDA encounters a decomposition step while solving a conjunctive query Q, it will end up partitioning a table R_T into a union of k *pairwise disjoint* tables $R_T^{(1)}, \ldots, R_T^{(k)}$, and accordingly the entire query Q is effectively partitioned into a union of k different queries $Q^{(1)}, \ldots, Q^{(k)}$, each of which contains one part $R_T^{(j)}$ of R_T. The analogous scenario in FAQ-SS queries is that a factor ψ_T is going to get partitioned into k factors $\psi_T^{(1)}, \ldots, \psi_T^{(k)}$ whose *supports* are pairwise disjoint. (The support of a factor ψ_S is the set of tuples $\mathbf{a}_S \in \prod_{i \in S} \mathsf{Dom}(A_i)$ such that $\psi_S(\mathbf{a}_S) \neq \mathbf{0}$ where $\mathbf{0}$ is the additive identity of the commutative semiring $(\mathbf{D}, \oplus, \otimes)$.) Therefore, the factor ψ_T can be written as a sum $\oplus$ of the k parts $\psi_T^{(1)}, \ldots, \psi_T^{(k)}$, i.e.

$$\psi_T(\mathbf{a}_T) = \bigoplus_{j \in [k]} \psi_T^{(j)}(\mathbf{a}_T).$$

As a result, the entire FAQ-SS query can be written as a sum of k FAQ-SS queries as follows.

$$\begin{aligned}
\varphi_F(\mathbf{a}_F) &= \bigoplus_{a_{f+1} \in \mathsf{Dom}(A_{f+1})} \cdots \bigoplus_{a_n \in \mathsf{Dom}(A_n)} \bigotimes_{S \in \mathcal{E}} \psi_S(\mathbf{a}_S) \\
&= \bigoplus_{a_{f+1} \in \mathsf{Dom}(A_{f+1})} \cdots \bigoplus_{a_n \in \mathsf{Dom}(A_n)} \left[\left(\bigoplus_{j \in [k]} \psi_T^{(j)}(\mathbf{a}_T) \right) \otimes \left(\bigotimes_{S \in \mathcal{E} \setminus \{T\}} \psi_S(\mathbf{a}_S) \right) \right] \\
&= \bigoplus_{a_{f+1} \in \mathsf{Dom}(A_{f+1})} \cdots \bigoplus_{a_n \in \mathsf{Dom}(A_n)} \bigoplus_{j \in [k]} \left[\psi_T^{(j)}(\mathbf{a}_T) \otimes \left(\bigotimes_{S \in \mathcal{E} \setminus \{T\}} \psi_S(\mathbf{a}_S) \right) \right] \qquad (75) \\
&= \bigoplus_{j \in [k]} \underbrace{\bigoplus_{a_{f+1} \in \mathsf{Dom}(A_{f+1})} \cdots \bigoplus_{a_n \in \mathsf{Dom}(A_n)} \left[\psi_T^{(j)}(\mathbf{a}_T) \otimes \left(\bigotimes_{S \in \mathcal{E} \setminus \{T\}} \psi_S(\mathbf{a}_S) \right) \right]}_{\varphi_F^{(j)}(\mathbf{a}_F)}. \qquad (76)
\end{aligned}$$

Corollary 6.15. *An* FAQ-SS *query Q without free variables and with degree constraints DC can be solved by* PANDA *in time*

$$\tilde{O}(N + \mathrm{poly}(\log N) \cdot 2^{\mathsf{da\text{-}subw}(Q)} + |\mathsf{output}|).$$

For general FAQ queries (where we have multiple types of aggregates), the above approach no longer works. In particular, the transition from (75) to (76) would fail if any two of the aggregates $\oplus^{(f+1)}, \ldots, \oplus^{(n)}$ were different operators.

7 Concluding remarks

Our negative answer to question 1 leads to a natural question: can we design an algorithm whose runtime matches the entropic bound under the presence of FDs or degree constraints? Worst-case optimal join algorithms [38, 39, 43] were able to achieve this when there is no FDs. And, as shown in [32] there are classes of queries with FDs for which the answer is positive (using the chain algorithm). A natural direction is to extend the class of queries with FDs that an algorithm meeting the entropic bound can be designed, beyond what was shown in [32]. Along the same line, the next natural open question is to design algorithms to evaluate disjunctive datalog queries matching the bound $\mathsf{LogSizeBound}_{\bar{\Gamma}_n^* \cap \mathsf{HDC}}(Q)$. From there, the posibility of achieving eda-subw and/or eda-fhtw is within reach. We already have an example where PANDA was able to achieve eda-subw and eda-fhtw: the 4-cycle example. Another big open question is to remove the polylog factor from the runtime of PANDA.

References

[1] Aji, S. M., and McEliece, R. J. The generalized distributive law. *IEEE Transactions on Information Theory 46*, 2 (2000), 325–343.

[2] Alon, N. On the number of subgraphs of prescribed type of graphs with a given number of edges. *Israel J. Math. 38*, 1-2 (1981), 116–130.

[3] Alon, N., Yuster, R., and Zwick, U. Finding and counting given length cycles. *Algorithmica 17*, 3 (1997), 209–223.

[4] Alviano, M., Faber, W., Leone, N., Perri, S., Pfeifer, G., and Terracina, G. The disjunctive datalog system DLV. In *Datalog Reloaded - First International Workshop, Datalog 2010, Oxford, UK, March 16-19, 2010. Revised Selected Papers* (2010), O. de Moor, G. Gottlob, T. Furche, and A. J. Sellers, Eds., vol. 6702 of *Lecture Notes in Computer Science*, Springer, pp. 282–301.

[5] Atserias, A., Grohe, M., and Marx, D. Size bounds and query plans for relational joins. In *FOCS* (2008), IEEE Computer Society, pp. 739–748.

[6] Bakibayev, N., Kociský, T., Olteanu, D., and Zavodny, J. Aggregation and ordering in factorised databases. *PVLDB 6*, 14 (2013), 1990–2001.

[7] Balister, P., and Bollobás, B. Projections, entropy and sumsets. *Combinatorica 32*, 2 (2012), 125–141.

[8] Benedikt, M., Leblay, J., and Tsamoura, E. Querying with access patterns and integrity constraints. *PVLDB 8*, 6 (2015), 690–701.

[9] Benedikt, M., ten Cate, B., and Tsamoura, E. Generating plans from proofs. *ACM Trans. Database Syst. 40*, 4 (2016), 22.

[10] Cao, Y., Fan, W., Wo, T., and Yu, W. Bounded conjunctive queries. *PVLDB 7*, 12 (2014), 1231–1242.

[11] Chan, T. H. Group characterizable entropy functions. *CoRR abs/cs/0702064* (2007).

[12] Chan, T. H., and Yeung, R. W. On a relation between information inequalities and group theory. *IEEE Transactions on Information Theory 48*, 7 (2002), 1992–1995.

[13] Chekuri, C., and Rajaraman, A. Conjunctive query containment revisited. *Theor. Comput. Sci. 239*, 2 (2000), 211–229.

[14] Chung, F. R. K., Graham, R. L., Frankl, P., and Shearer, J. B. Some intersection theorems for ordered sets and graphs. *J. Combin. Theory Ser. A 43*, 1 (1986), 23–37.

[15] Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. *Introduction to Algorithms, Third Edition*, 3rd ed. The MIT Press, 2009.

[16] Dechter, R. Bucket elimination: A unifying framework for reasoning. *Artif. Intell. 113*, 1-2 (1999), 41–85.

[17] Du, D., and Pardalos, P. *Minimax and Applications*. Nonconvex Optimization and Its Applications. Springer US, 1995.

[18] Edmonds, J., and Karp, R. M. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM 19*, 2 (1972), 248–264.

[19] Eiter, T., Gottlob, G., and Mannila, H. Disjunctive datalog. *ACM Trans. Database Syst. 22*, 3 (1997), 364–418.

[20] Fagin, R. Degrees of acyclicity for hypergraphs and relational database schemes. *J. ACM 30*, 3 (1983), 514–550.

[21] Fischl, W., Gottlob, G., and Pichler, R. General and Fractional Hypertree Decompositions: Hard and Easy Cases. *ArXiv e-prints* (Nov. 2016).

[22] Flum, J., Frick, M., and Grohe, M. Query evaluation via tree-decompositions. *J. ACM 49*, 6 (2002), 716–752.

[23] Freuder, E. C. Complexity of k-tree structured constraint satisfaction problems. In *Proceedings of the 8th National Conference on Artificial Intelligence. Boston, Massachusetts, July 29 - August 3, 1990, 2 Volumes.* (1990), H. E. Shrobe, T. G. Dietterich, and W. R. Swartout, Eds., AAAI Press / The MIT Press, pp. 4–9.

[24] Friedgut, E., and Kahn, J. On the number of copies of one hypergraph in another. *Israel J. Math. 105* (1998), 251–256.

[25] Gogacz, T., and Toruńczyk, S. Entropy bounds for conjunctive queries with functional dependencies. In *Proc. 20th International Conference on Database Theory (ICDT)* (2017). To appear.

[26] Gottlob, G., Greco, G., Leone, N., and Scarcello, F. Hypertree decompositions: Questions and answers. In Milo and Tan [37], pp. 57–74.

[27] Gottlob, G., Lee, S. T., Valiant, G., and Valiant, P. Size and treewidth bounds for conjunctive queries. *J. ACM 59*, 3 (2012), 16.

[28] Gottlob, G., Leone, N., and Scarcello, F. Hypertree decompositions and tractable queries. *J. Comput. Syst. Sci. 64*, 3 (2002), 579–627.

[29] Grohe, M., and Marx, D. Constraint solving via fractional edge covers. In *SODA* (2006), ACM Press, pp. 289–298.

[30] Grohe, M., and Marx, D. Constraint solving via fractional edge covers. *ACM Transactions on Algorithms 11*, 1 (2014), 4.

[31] Khamis, M. A., Ngo, H. Q., and Rudra, A. FAQ: questions asked frequently. In Milo and Tan [37], pp. 13–28.

[32] Khamis, M. A., Ngo, H. Q., and Suciu, D. Computing join queries with functional dependencies. In Milo and Tan [37], pp. 327–342.

[33] Koch, C. Incremental query evaluation in a ring of databases. In *Proceedings of the Twenty-ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (New York, NY, USA, 2010), PODS '10, ACM, pp. 87–98.

[34] Marx, D. Tractable structures for constraint satisfaction with truth tables. *Theory Comput. Syst. 48*, 3 (2011), 444–464.

[35] Marx, D. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM 60*, 6 (2013), Art. 42, 51.

[36] Matus, F. Infinitely many information inequalities. In *2007 IEEE International Symposium on Information Theory* (2007), IEEE, pp. 41–44.

[37] Milo, T., and Tan, W., Eds. *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016* (2016), ACM.

[38] Ngo, H. Q., Porat, E., Ré, C., and Rudra, A. Worst-case optimal join algorithms: [extended abstract]. In *PODS* (2012), pp. 37–48.

[39] Ngo, H. Q., Ré, C., and Rudra, A. Skew strikes back: new developments in the theory of join algorithms. *SIGMOD Record 42*, 4 (2013), 5–16.

[40] Schrijver, A. *Theory of linear and integer programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons Ltd., Chichester, 1986. A Wiley-Interscience Publication.

[41] Tarjan, R. E., and Yannakakis, M. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM J. Comput. 13*, 3 (July 1984), 566–579.

[42] Vardi, M. Y. The complexity of relational query languages (extended abstract). In *STOC* (1982), pp. 137–146.

[43] Veldhuizen, T. L. Triejoin: A simple, worst-case optimal join algorithm. In *Proc. 17th International Conference on Database Theory (ICDT), Athens, Greece, March 24-28, 2014.* (2014), N. Schweikardt, V. Christophides, and V. Leroy, Eds., OpenProceedings.org, pp. 96–106.

[44] Yannakakis, M. Algorithms for acyclic database schemes. In *VLDB* (1981), pp. 82–94.

[45] Yeung, R. W. *Information Theory and Network Coding*, 1 ed. Springer Publishing Company, Incorporated, 2008.

[46] Zhang, Z., and Yeung, R. W. On characterization of entropy function via information inequalities. *IEEE Transactions on Information Theory 44*, 4 (1998), 1440–1452.

A More details on Section 2

Proof of Proposition 2.7. Let $(\lambda_F)_{F\in\mathcal{E}}$ be any integral edge cover of $\mathcal{H}$. Let $\{F_1,\ldots,F_k\}$ be the collection of hyperedges in $\mathcal{E}$ such that $\lambda_{F_i}=1, \forall i\in[k]$. Define $F_0=\emptyset$. Then from subadditivity, monotonicity, and the cardinality constraints we have

$$\begin{aligned}
& & h(V) & \\
\text{(sub-additivity)} & \leq & \sum_{i=1}^{k} h\left(F_i \setminus \bigcup_{j=1}^{i-1} F_j\right) & \\
\text{(monotonicity)} & \leq & \sum_{i=1}^{k} h\left(F_i\right) & \\
& = & \sum_{F\in\mathcal{E}} \lambda_F \cdot h(F) & \\
\text{(cardinality constraints)} & \leq & \sum_{F\in\mathcal{E}} \lambda_F \log N_F & \\
& = & \log_2 \left(\prod_{F\in\mathcal{E}} N_F^{\lambda_F}\right). &
\end{aligned}$$

Conversely, suppose h^* is an optimal solution to the linear program

$$\begin{aligned} \max \quad & h(V) && \\ \text{s.t.} \quad & h(I \cup J) \le h(I) + h(J), && I \perp J \\ & h(F) \le \log N_F, && F \in \mathcal{E} \\ & h(X) \le h(Y), && X \subset Y \subseteq V \\ & h(Z) \ge 0 && \forall \emptyset \neq Z \subseteq V. \end{aligned} \tag{77}$$

We show that there is an integral edge cover $(\lambda_F)_{F\in\mathcal{E}}$ such that $\sum_{F\in\mathcal{E}} \lambda_F \log_2 N_F \le h^*(V)$. We need the dual of the LP (77). Associate a dual variable δ_F to each cardinality constraint $h(F) \le \log N_F$, a dual variable $\sigma_{I,J}$ to each sub-additivity constraint, and a dual variable $\mu_{X,Y}$ to each monotonicity constraint. For convenience, let $\delta_Z \stackrel{\text{def}}{=} 0$ for any $Z \in 2^V \setminus \mathcal{E}$. For any $Z \subseteq V$, define $\mathsf{flow}(Z)$ to be

$$\mathsf{flow}(Z) \stackrel{\text{def}}{=} \delta_Z + \sum_{I\perp J : I\cup J = Z} \sigma_{I,J} - \sum_{J : J\perp Z} \sigma_{Z,J} - \sum_{X : X\subset Z} \mu_{X,Z} + \sum_{Y : Z\subset Y} \mu_{Z,Y}. \tag{78}$$

Then, the dual LP is

$$\begin{aligned} \min \quad & \textstyle\sum_{F\in\mathcal{E}} (\log N_F)\cdot \delta_F && \\ \text{s.t.} \quad & \mathsf{flow}(Z) && \ge 0, \quad \emptyset \neq Z \subseteq V \\ & \mathsf{flow}(V) && \ge 1 \\ & (\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu}) && \ge \mathbf{0}. \end{aligned} \tag{79}$$

Claim 1: Given a feasible dual solution $(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$, let D be a common denominator for all entries in $(\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$. Then, there exist D integral edge covers $(\lambda_F^{(1)})_{F\in\mathcal{E}}, (\lambda_F^{(2)})_{F\in\mathcal{E}}, \ldots, (\lambda_F^{(D)})_{F\in\mathcal{E}}$ such that for any $F \in \mathcal{E}$, $\sum_{i=1}^{D} \lambda_F^{(i)} \le D \cdot \delta_F$.

Assuming Claim 1 holds, the proof works as follows. Let $(\boldsymbol{\delta}^*, \boldsymbol{\sigma}^*, \boldsymbol{\mu}^*)$ be an optimal dual solution, then $h^*(V) = \sum_{F\in\mathcal{E}} (\log N_F)\cdot \delta_F^*$ by strong duality. Applying Claim 1 on $(\boldsymbol{\delta}^*, \boldsymbol{\sigma}^*, \boldsymbol{\mu}^*)$, we get

$$h^*(V) \ge \sum_{F\in\mathcal{E}} \sum_{i=1}^{D} \frac{\lambda_F^{(i)}}{D} \log N_F \ge \min_{i\in[D]} \sum_{F\in\mathcal{E}} \lambda_F^{(i)} \log N_F.$$

Proof of Claim 1: Let $w \stackrel{\text{def}}{=} 1/D$. Initially, let $\mathcal{B} = \{V\}$. If $\mathsf{flow}(B) > \delta_B$ for some $B \neq \emptyset$, then from (78) there must be either (1) $I \perp J$ with $I \cup J = B$ and $\sigma_{I,J} > 0$, or (2) $Y \supset B$ with $\mu_{B,Y} > 0$. While there is $B \in \mathcal{B}$ with $\mathsf{flow}(B) > \delta_B$, we do the following:

- If Case (1) holds, then we reduce $\sigma_{I,J}$ by w, remove B from $\mathcal{B}$, and add both I and J to $\mathcal{B}$.
- If Case (2) holds, then we reduce $\mu_{B,Y}$ by w, remove B from $\mathcal{B}$, and add Y to $\mathcal{B}$.

Either way, we are maintaining the following invariants (which initially hold):

- $\mathsf{flow}(B) > 0$ for any $B \in \mathcal{B}$.
- $\mathsf{flow}(Z) \ge 0$ for any $\emptyset \neq Z \subseteq V$.
- For any $v \in V$, there is some $B \in \mathcal{B}$ such that $v \in B$.

The above process terminates because at each iteration we are reducing $\|\boldsymbol{\sigma}\|_1 + \|\boldsymbol{\mu}\|_1$ by w. When it does terminate, we have $\delta_B \ge \mathsf{flow}(B) > 0$ for all $B \in \mathcal{B}$. Hence, $(\lambda_F^{(1)})_{F\in\mathcal{E}}$ is an integral edge cover where

$$\lambda_F^{(1)} \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } F \in \mathcal{B}, \\ 0 & \text{otherwise.} \end{cases}$$

Now for every $B \in \mathcal{B}$, we reduce δ_B (hence $\mathsf{flow}(B)$) by w and remove B from $\mathcal{B}$. Notice that except for the initial V in $\mathcal{B}$, the following was always maintained. Before we added any B to $\mathcal{B}$ we increased $\mathsf{flow}(B)$ by w, and before we removed any B' from $\mathcal{B}$ we decreased $\mathsf{flow}(B')$ by w. Therefore, at the end $\mathsf{flow}(Z)$ is unchanged for all $Z \subset V$, and $\mathsf{flow}(V)$ is reduced by w. If the final $\mathsf{flow}(V) = 0$, then we are done. Otherwise, $\frac{1}{1-w} \cdot (\boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ is a feasible dual solution with $(D-1)$ as a common denominator, and we can repeat the process. □

Proof of Proposition 2.3. Every non-dominated non-redundant tree decomposition can be constructed from a variable ordering [31], and there are $n!$ variable orderings. When we run GYO-elimination on such a tree decomposition with an arbitrary root-bag, every bag contains a variable that does not belong to the parent bag. □

B More details on Section 4

B.1 Missing Proofs

Proof of Proposition 4.3. The proposition is a simple consequence of Farkas' lemma. There are *many* variants of Farkas' lemma [40]. We use a version whose proof we also reproduce here because the proof is very short.

Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be a matrix and $\mathbf{c} \in \mathbb{R}^n$ be a vector. Let $P = \{\mathbf{x} \mid \mathbf{A}\mathbf{x} \le \mathbf{0}, \mathbf{x} \ge \mathbf{0}\}$ be a polyhedron and $D = \{\mathbf{y} \mid \mathbf{A}^T\mathbf{y} \ge \mathbf{c}, \mathbf{y} \ge \mathbf{0}\}$ be a dual polyhedron. Then, a variant of Farkas' lemma states that D is non-empty if and only if there is *no* $\mathbf{x} \in P$ such that $\mathbf{c}^T\mathbf{x} > 0$. To see this, note that the system $\{\mathbf{c}^T\mathbf{x} > 0, \mathbf{x} \in P\}$ is infeasible iff $\max\{\mathbf{c}^T\mathbf{x} \mid \mathbf{x} \in P\} = 0$, which by strong duality is equivalent to $\min\{\mathbf{0}^T\mathbf{y} \mid \mathbf{y} \in D\}$ is feasible, which is the same as D is non-empty.

Now, to see why the above variant of Farkas' lemma implies Proposition 4.3, we note that (43) holds for all polymatroids iff $\{(\boldsymbol{\lambda}_{\mathcal{B}} - \boldsymbol{\delta}_{\mathsf{DC}})^T\mathbf{h} > 0 \mid \mathbf{h} \in \Gamma_n\}$ is infeasible; now we are in the exact setting of the above variant of Farkas' lemma and the rest follows trivially. □

Proof of Lemma 4.2. The dual of (40) can be written as

$$\begin{aligned} \min \quad & \mathbf{b}^T\mathbf{y} \\ \text{s.t.} \quad & \mathbf{A}^T\mathbf{y} \ge \mathbf{C}\boldsymbol{\lambda} \\ & \mathbf{y} \ge \mathbf{0}. \end{aligned} \tag{80}$$

Let $\mathbf{1}_p \in \mathbb{R}^p$ be the all-1 vector. The LP (39) can be rewritten as

$$\begin{aligned} \max \quad & w \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} \le \mathbf{b} \\ & \mathbf{1}_p w - \mathbf{C}^T\mathbf{x} \le \mathbf{0} \\ & \mathbf{x} \ge \mathbf{0}, \quad w \ge 0. \end{aligned} \tag{81}$$

The dual of (81) is

$$\begin{aligned} \min \quad & \mathbf{b}^T\mathbf{y} \\ \text{s.t.} \quad & \mathbf{A}^T\mathbf{y} \ge \mathbf{C}\mathbf{z} \\ & \mathbf{1}_p^T\mathbf{z} \ge 1 \\ & \mathbf{y} \ge \mathbf{0}, \quad \mathbf{z} \ge 0. \end{aligned} \tag{82}$$

Let $(w^*, \mathbf{x}^*)$ and $(\mathbf{z}^*, \mathbf{y}^*)$ be a pair of primal-optimal and dual-optimal solutions to (81) and (82). Define $\boldsymbol{\lambda} = \mathbf{z}^*$. To show that (39) is equivalent to (40), we show that (81) is equivalent to (40). In particular, we claim that $\mathbf{x}^*$ and $\mathbf{y}^*$ are a pair of primal-optimal and dual-optimal solution to (81) and (80), and that the objective values of (40) and (81) are identical.

The fact that $\mathbf{x}^*$ and $\mathbf{y}^*$ are feasible to (40) and (80) is trivial. We are left to verify that they have the same objective value. Due to complementary slackness of the (81) and (82) pair, note that $w^* > 0$ implies $\mathbf{1}_p^T \mathbf{z}^* = 1$ and $(\mathbf{1}_p^T w^* - \mathbf{C}^T \mathbf{x}^*)^T \mathbf{z}^* = 0$. It follows that

$$\mathbf{b}^T \mathbf{y}^* = w^* = (\mathbf{1}_p^T \mathbf{z}^*) \cdot w^* = (\mathbf{C}^T \mathbf{x}^*)^T \cdot \mathbf{z}^* = (\mathbf{C}\mathbf{z}^*)^T \cdot \mathbf{x}^* = (\mathbf{C}\boldsymbol{\lambda})^T \cdot \mathbf{x}^*.$$

□

B.2 Bounding $\|\boldsymbol{\delta}\|_1, \|\boldsymbol{\mu}\|_1$, and $\|\boldsymbol{\sigma}\|_1$ (w.r.t $\|\boldsymbol{\lambda}\|_1$)

Given a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with a witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, there are various ways to construct a proof sequence for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$. (Theorem 4.9 showed one possible construction, and we will see more efficient constructions in the next section.) What those various constructions have in common is that the length of the resulting proof sequence depends on $\|\boldsymbol{\sigma}\|_1$, $\|\boldsymbol{\delta}\|_1$ and $\|\boldsymbol{\mu}\|_1$. In turn, the runtime of PANDA algorithm depends on the length of the proof sequence, which provides a motivation for minimizing $\|\boldsymbol{\sigma}\|_1$, $\|\boldsymbol{\delta}\|_1$ and $\|\boldsymbol{\mu}\|_1$ as much as possible, which is our target in this section. In particular, we want to replace the original inequality (and its witness) with a new inequality that is "just as good" but has lower $\|\boldsymbol{\sigma}\|_1$, $\|\boldsymbol{\delta}\|_1$ and $\|\boldsymbol{\mu}\|_1$.

This section is outlined as follows. We start Section B.2.1 by providing some definitions that formalize what we meant by an inequality being "just as good" as another, then we proceed to bounding the total of monotonicity terms of the form $\mu_{X,Y}$ for $X \neq \emptyset$. In Section B.2.2, we bound the total of terms of the form $\delta_{Y|\emptyset}$. Section B.2.3 shows that this is equivalent to bounding monotonicity terms $\mu_{\emptyset,Y}$, hence we now have a bound on all monotonicity terms $\|\boldsymbol{\mu}\|_1$. Finally, Section B.2.4 shows that bounding $\|\boldsymbol{\mu}\|_1$ imposes a bound on both $\|\boldsymbol{\delta}\|_1$ and $\|\boldsymbol{\sigma}\|_1$.

B.2.1 Bounding the total $\mu_{X,Y}$ for $X \neq \emptyset$

Definition B.1 ($\boldsymbol{\delta}'$ dominated by $\boldsymbol{\delta}$)**.** Given two vectors $\boldsymbol{\delta}, \boldsymbol{\delta}' \in \mathbb{Q}_+^{\mathcal{P}}$, we say that $\boldsymbol{\delta}'$ is *dominated* by $\boldsymbol{\delta}$ if there exists a function $f : \mathcal{P} \times \mathcal{P} \to \mathbb{Q}_+$ that satisfies the following conditions.

- If $f((X,Y),(X',Y')) \neq 0$, then $Y' \setminus X' \subseteq Y \setminus X$ and $X' \supseteq X$.
- $\forall (X,Y) \in \mathcal{P}, \quad \sum_{(X',Y') \in \mathcal{P}} f((X,Y),(X',Y')) \le \delta_{Y|X}$.
- $\forall (X',Y') \in \mathcal{P}, \quad \sum_{(X,Y) \in \mathcal{P}} f((X,Y),(X',Y')) \ge \delta'_{Y'|X'}$.

The function f is called the *domination function for* $(\boldsymbol{\delta}, \boldsymbol{\delta}')$.

Definition B.2 (Rational compatibility)**.** Given two vectors of rational numbers $\mathbf{q} \in \mathbb{Q}_+^l, \mathbf{q}' \in \mathbb{Q}_+^{l'}$(for some integers l, l'), we say that $\mathbf{q}'$ is *rationally-compatible* with $\mathbf{q}$ if the minimum common denominator of all entries in $\mathbf{q}$ is a common denominator (not-necessarily minimum) for all entries in $\mathbf{q}'$.

Lemma B.3. *For any Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with a witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, there exists a Shannon flow inequality $\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ with a tight witness $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ such that $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$ and $\boldsymbol{\lambda}'$ dominates $\boldsymbol{\lambda}$ and $\mu'_{X,Y} = 0$ for all $\emptyset \neq X \subset Y \subseteq V$. Moreover, $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$ is rationally-compatible with $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$.*

Proof. W.L.O.G we can assume the witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$ to be tight. We start by choosing $\boldsymbol{\delta}' = \boldsymbol{\delta}$, $\boldsymbol{\lambda}' = \boldsymbol{\lambda}$, $(\boldsymbol{\sigma}', \boldsymbol{\mu}') = (\boldsymbol{\sigma}, \boldsymbol{\mu})$. Moreover, we choose the domination function f for $(\boldsymbol{\delta}, \boldsymbol{\delta}')$ to be

$$f((X,Y),(X',Y')) = \begin{cases} \delta_{Y|X} & \text{if } X' = X \text{ and } Y' = Y \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, the domination function g for $(\boldsymbol{\lambda}', \boldsymbol{\lambda})$ is chosen to be

$$g((X',Y'),(X,Y)) = \begin{cases} \lambda_{Y|X} & \text{if } X' = X \text{ and } Y' = Y \\ 0 & \text{otherwise.} \end{cases}$$

Let D be the common dominator of $(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$, and w be $1/D$. While there are $\emptyset \neq X \subset Y \subseteq V$ where $\mu'_{X,Y} > 0$ we apply the following. If $\lambda'_X > 0$, then we reduce both λ'_X and $\mu'_{X,Y}$ by w and increase λ'_Y by w. Moreover, we choose an arbitrary Z (if any) where $g((\emptyset, X), (\emptyset, Z)) > 0$, and we reduce $g((\emptyset, X), (\emptyset, Z))$ by w and increase $g((\emptyset, Y), (\emptyset, Z))$ by w. Otherwise (if $\lambda'_X = 0$), since the witness $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ is tight and $\mu'_{X,Y}$ is increasing $\mathsf{inflow}(X)$, there must be some dual variable that is reducing $\mathsf{inflow}(X)$ back to 0. In particular, we recognize three cases: (See Figure 3 [right].)

1. If there is $W \subset X$ such that $\mu'_{W,X} > 0$, then we reduce both $\mu'_{W,X}$ and $\mu'_{X,Y}$ by w and increase $\mu'_{W,Y}$ by w.
2. If there is some $Y' \supset X$ such that $\delta'_{Y'|X} > 0$, then we reduce both $\mu'_{X,Y}$ and $\delta'_{Y'|X}$ by w and increase both $\delta'_{Y \cup Y'|Y}$ and $\mu'_{Y', Y \cup Y'}$ by w. Moreover, we choose an arbitrary $W \subset Z \subseteq V$ (if any) where $f((W, Z), (X, Y')) > 0$, and we reduce $f((W, Z), (X, Y'))$ by w and increase $f((W, Z), (Y, Y \cup Y'))$ by w.
3. If there is some $X' \perp X$ such that $\sigma'_{X,X'} > 0$, then we reduce both $\mu'_{X,Y}$ and $\sigma'_{X,X'}$ by w and increase each one of $\sigma'_{Y,X'}$, $\mu'_{X \cup X', Y \cup X'}$, $\mu'_{X \cap X', Y \cap X'}$ by w.

In all cases, we are maintaining $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ as a valid and tight witness. We are also maintaining that $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$ and $\boldsymbol{\lambda}'$ dominates $\boldsymbol{\lambda}$. To prove that the above process terminates, consider the following non-negative, bounded, integral function.

$$\phi(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}') \stackrel{\text{def}}{=} D \cdot \left[\sum_B \lambda'_B (n - |B|) + \sum_{X \subset Y} \delta'_{Y|X} (|Y| - |X|)(n - |X|) + \sum_{I \perp J} \sigma'_{I,J} (2n - |I| - |J|) \right].$$

In all cases except (1), $\phi(\boldsymbol{\lambda}', \boldsymbol{\delta}', \boldsymbol{\sigma}')$ decreases by 1. Case (1) reduces the quantity $D \cdot \|\boldsymbol{\mu}'\|_1$ by 1. Initially, we have $\|\boldsymbol{\mu}'\|_1 = \|\boldsymbol{\mu}\|_1$. However, $\|\boldsymbol{\mu}'\|_1$ increases in case (3). Hence, initially case (1) can maximally be repeated $D \cdot \|\boldsymbol{\mu}\|_1$ consecutive times. Later on, every time we increase some $\mu'_{X,Y}$ in cases (2) and (3), we can immediately check whether there is some $W \subset X$ where $\mu'_{W,X} > 0$ and if so apply case (1). By doing so, we can amortize the cost of case (1) over cases (2) and (3). □

Corollary B.4 (The total $\mu_{X,Y}$ for $X \neq \emptyset$ is $\leq \|\boldsymbol{\lambda}\|_1$)**.** *For any Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with a witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, there exists a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ with a tight witness $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ such that $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$ and*

$$\forall \emptyset \neq X \subseteq V, \quad \sum_{Y \supset X} \mu'_{X,Y} \leq \lambda_X.$$

Hence

$$\sum_{\emptyset \neq X \subset Y \subseteq V} \mu'_{X,Y} \leq \|\boldsymbol{\lambda}\|_1. \tag{83}$$

Moreover, $(\boldsymbol{\lambda}, \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$ is rationally-compatible with $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$.

B.2.2 Bounding the total $\delta_{Y|\emptyset}$

Lemma B.5 (The total $\delta_{Y|\emptyset}$ is $\leq n \cdot \|\boldsymbol{\lambda}\|_1$)**.** *For any Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with a witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, there exists a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ with a witness $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ such that $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$ and*

$$\forall v \in V, \quad \sum_{Y \ni v} \delta'_{Y|\emptyset} \leq \|\boldsymbol{\lambda}\|_1.$$

Hence

$$\sum_{Y \subseteq V} \delta'_{Y|\emptyset} \leq n \cdot \|\boldsymbol{\lambda}\|_1. \tag{84}$$

Moreover, $(\boldsymbol{\lambda}, \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$ is rationally-compatible with $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$.

Proof. Fix an arbitrary variable $v \in V$ such that $\sum_{Y \ni v} \delta_{Y|\emptyset} > \|\boldsymbol{\lambda}\|_1$. Define a function $\mathcal{F} : \mathbb{Q}_+^{\mathcal{P}} \to \mathbb{Q}_+^{\mathcal{P}}$ such that for any vector $\mathbf{t} \in \mathbb{Q}_+^{\mathcal{P}}$, $\mathcal{F}(\mathbf{t})$ is a vector $\mathbf{t}' \in \mathbb{Q}_+^{\mathcal{P}}$ defined as

$$t'_{Y|X} \stackrel{\text{def}}{=} \begin{cases} t_{Y|X} + t_{Y|X-\{v\}} + t_{Y-\{v\}|X-\{v\}} & \text{if } v \in X \text{ (hence } v \in Y) \\ 0 & \text{otherwise (i.e. if } v \notin Y \text{ or } v \notin X) \end{cases}$$

Note that $\mathbf{t}'$ is dominated by $\mathbf{t}$ where the domination function f is defined as

$$\begin{aligned} f((X,Y),(X \cup \{v\}, Y \cup \{v\})) &= t_{Y|X}, \\ f((X,Y),(X',Y')) &= 0 \qquad \text{otherwise.} \end{aligned}$$

Claim 1: The inequality $\langle \mathcal{F}(\boldsymbol{\lambda}), \mathbf{h} \rangle \le \langle \mathcal{F}(\boldsymbol{\delta}), \mathbf{h} \rangle$ holds for all conditional polymatroids $\mathbf{h}$.

(Notice that the inequality $\langle \mathcal{F}(\boldsymbol{\lambda}), \mathbf{h} \rangle \le \langle \mathcal{F}(\boldsymbol{\delta}), \mathbf{h} \rangle$ does not necessarily have the form (37) of a Shannon flow inequality, because it can have $\lambda'_{B \cup \{v\}|\{v\}} > 0$.) We will prove Claim 1 based on the following claim.

Claim 2: For any $\mathbf{f} \in \{\mathbf{s}_{I,J}, \mathbf{m}_{X,Y}, \mathbf{c}_{X,Y}, \mathbf{d}_{Y,X}\}$ and any conditional polymatroid $\mathbf{h}$, the following inequality holds

$$\langle \mathcal{F}(\mathbf{f}), \mathbf{h} \rangle \le 0.$$

Claim 2 holds because the following inequalities hold for all conditional polymatroids.

$$\begin{aligned} h(I \cup J \cup \{v\}|J \cup \{v\}) - h(I \cup \{v\}|(I \cap J) \cup \{v\}) &\le 0, I \perp J \\ -h(Y \cup \{v\}|\{v\}) + h(X \cup \{v\}|\{v\}) &\le 0, X \subset Y \\ h(Y \cup \{v\}|\{v\}) - h(Y \cup \{v\}|X \cup \{v\}) - h(X \cup \{v\}|\{v\}) &\le 0, X \subset Y \\ -h(Y \cup \{v\}|\{v\}) + h(Y \cup \{v\}|X \cup \{v\}) + h(X \cup \{v\}|\{v\}) &\le 0, X \subset Y \end{aligned}$$

To prove Claim 1, we construct a proof sequence for the original inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$, which is possible thanks to Theorem 4.9. Consider the corresponding inequality sequence $\boldsymbol{\delta} = \boldsymbol{\delta}_0, \boldsymbol{\delta}_1, \ldots, \boldsymbol{\delta}_\ell \ge \boldsymbol{\lambda}$ such that, for every $i \in [\ell]$, $\boldsymbol{\delta}_i = \boldsymbol{\delta}_{i-1} + w_i \cdot \mathbf{f}_i$ for some $w_i > 0$ and $\mathbf{f}_i \in \{\mathbf{s}_{I,J}, \mathbf{m}_{X,Y}, \mathbf{c}_{X,Y}, \mathbf{d}_{Y,X}\}$.

$$\begin{aligned} \langle \mathcal{F}(\mathbf{f}_i), \mathbf{h} \rangle &\le 0 \qquad \text{(Claim 2)} \\ \langle \mathcal{F}(1/w_i(\boldsymbol{\delta}_i - \boldsymbol{\delta}_{i-1})), \mathbf{h} \rangle &\le 0 \\ \langle 1/w_i(\mathcal{F}(\boldsymbol{\delta}_i) - \mathcal{F}(\boldsymbol{\delta}_{i-1})), \mathbf{h} \rangle &\le 0 \qquad \text{(by linearity of } \mathcal{F}) \\ \langle \mathcal{F}(\boldsymbol{\delta}_i), \mathbf{h} \rangle &\le \langle \mathcal{F}(\boldsymbol{\delta}_{i-1}), \mathbf{h} \rangle \end{aligned}$$

which proves Claim 1.

Finally, let $\boldsymbol{\delta}' = \mathcal{F}(\boldsymbol{\delta})$, $\boldsymbol{\lambda}' = \mathcal{F}(\boldsymbol{\lambda})$, f be the domination function of $(\boldsymbol{\delta}, \boldsymbol{\delta}')$, D be the common denominator of $(\boldsymbol{\delta}, \boldsymbol{\delta}', \boldsymbol{\lambda}, \boldsymbol{\lambda}')$, and $w = 1/D$. Initially, we have $\sum_{Y \ni v} \delta'_{Y|\emptyset} = 0$. While there is $B \subseteq V$ such that $\lambda'_{B \cup \{v\}|\{v\}} > 0$, finds $Y \ni v$ such that $f((\emptyset, Y), (\{v\}, Y)) > 0$, and reduce $\lambda'_{B \cup \{v\}|\{v\}}$, $f((\emptyset, Y), (\{v\}, Y))$, and $\delta'_{Y|\{v\}}$ by w, and increase $\lambda'_{B \cup \{v\}|\emptyset}$, $f((\emptyset, Y), (\emptyset, Y))$, and $\delta'_{Y|\emptyset}$ by w. Note that this update maintains that $\langle \boldsymbol{\lambda}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ holds for all polymatroids $\mathbf{h}$ and that $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$. At the end, $\boldsymbol{\lambda}'$ will dominate $\boldsymbol{\lambda}$ and we will have $\sum_{Y \ni v} \delta'_{Y|\emptyset} = \|\boldsymbol{\lambda}\|_1$ □

B.2.3 Bounding $\|\boldsymbol{\mu}\|_1$

Corollary B.6 ($\|\boldsymbol{\mu}\|_1 \le n \cdot \|\boldsymbol{\lambda}\|_1$)**.** *For any Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with a witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, there exists a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}'', \mathbf{h} \rangle$ with a tight witness $(\boldsymbol{\sigma}'', \boldsymbol{\mu}'')$ such that $\boldsymbol{\delta}''$ is dominated by $\boldsymbol{\delta}$ and*

$$\|\boldsymbol{\mu}''\|_1 \le n \cdot \|\boldsymbol{\lambda}\|_1. \tag{85}$$

In particular

$$\forall \emptyset \ne X \subseteq V, \quad \sum_{Y \supset X} \mu''_{X,Y} \le \lambda_X, \quad \textit{hence} \quad \sum_{\emptyset \ne X \subset Y \subseteq V} \mu''_{X,Y} \le \|\boldsymbol{\lambda}\|_1 \tag{86}$$

and

$$\sum_{Y\subseteq V} \mu''_{\emptyset,Y} \le (n-1)\cdot\|\boldsymbol{\lambda}\|_1. \tag{87}$$

Moreover, $(\boldsymbol{\lambda},\boldsymbol{\delta}'',\boldsymbol{\sigma}'',\boldsymbol{\mu}'')$ *is rationally-compatible with* $(\boldsymbol{\lambda},\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu})$.

Proof. First, we will apply Lemma B.5 to get a Shannon flow inequality $\langle\boldsymbol{\lambda},\mathbf{h}\rangle \le \langle\boldsymbol{\delta}',\mathbf{h}\rangle$ where $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$ and $\sum_{Y\subseteq V}\delta'_{Y|\emptyset} \le n\cdot\|\boldsymbol{\lambda}\|_1$. Then, we apply Corollary B.4 on $\langle\boldsymbol{\lambda},\mathbf{h}\rangle \le \langle\boldsymbol{\delta}',\mathbf{h}\rangle$ to get another Shannon flow inequality $\langle\boldsymbol{\lambda},\mathbf{h}\rangle \le \langle\boldsymbol{\delta}'',\mathbf{h}\rangle$ with a tight witness $(\boldsymbol{\sigma}'',\boldsymbol{\mu}'')$ such that $\boldsymbol{\delta}''$ is dominated by $\boldsymbol{\delta}'$ and (86) holds. Because $\boldsymbol{\delta}''$ is dominated by $\boldsymbol{\delta}'$, we have

$$\sum_{Y\subseteq V}\delta''_{Y|\emptyset} \le \sum_{Y\subseteq V}\delta'_{Y|\emptyset} \le n\cdot\|\boldsymbol{\lambda}\|_1.$$

Let $\mathsf{inflow}''(Z)$ denote the quantity $\mathsf{inflow}(Z)$ measured on the vector $(\boldsymbol{\delta}'',\boldsymbol{\sigma}'',\boldsymbol{\mu}'')$.

$$\begin{aligned}\sum_{\emptyset\neq Z\subseteq V}\mathsf{inflow}''(Z) &= \|\boldsymbol{\lambda}\|_1\\ \sum_{Z\neq\emptyset}\delta''_{Z|\emptyset} - \sum_{Z\neq\emptyset}\mu''_{\emptyset,Z} - \sum_{\substack{I\perp J\\ I\cap J=\emptyset}}\sigma''_{I,J} &= \|\boldsymbol{\lambda}\|_1\end{aligned}$$

$$\sum_{Z\neq\emptyset}\mu''_{\emptyset,Z} \le \sum_{Z\neq\emptyset}\delta''_{Z|\emptyset} - \|\boldsymbol{\lambda}\|_1 \le n\cdot\|\boldsymbol{\lambda}\|_1 - \|\boldsymbol{\lambda}\|_1.$$

$$\|\boldsymbol{\mu}''\|_1 = \sum_{Z\neq\emptyset}\mu''_{\emptyset,Z} + \sum_{\emptyset\neq X\subset Y\subseteq V}\mu''_{X,Y} \le n\cdot\|\boldsymbol{\lambda}\|_1.$$

□

B.2.4 Bounding $\|\boldsymbol{\sigma}\|_1$ and $\|\boldsymbol{\delta}\|_1$

Corollary B.7 $(2\|\boldsymbol{\sigma}\|_1+\|\boldsymbol{\delta}\|_1 \le n^3\cdot\|\boldsymbol{\lambda}\|_1)$**.** *For any Shannon flow inequality* $\langle\boldsymbol{\lambda},\mathbf{h}\rangle \le \langle\boldsymbol{\delta},\mathbf{h}\rangle$ *with a witness* $(\boldsymbol{\sigma},\boldsymbol{\mu})$*, there exists a Shannon flow inequality* $\langle\boldsymbol{\lambda},\mathbf{h}\rangle \le \langle\boldsymbol{\delta}'',\mathbf{h}\rangle$ *with a tight witness* $(\boldsymbol{\sigma}'',\boldsymbol{\mu}'')$ *such that* $\boldsymbol{\delta}''$ *is dominated by* $\boldsymbol{\delta}$ *and*

$$2\|\boldsymbol{\sigma}''\|_1+\|\boldsymbol{\delta}''\|_1 \le n^3\cdot\|\boldsymbol{\lambda}\|_1, \tag{88}$$

$$\|\boldsymbol{\mu}''\|_1 \le n\cdot\|\boldsymbol{\lambda}\|_1. \tag{89}$$

Moreover, $(\boldsymbol{\lambda},\boldsymbol{\delta}'',\boldsymbol{\sigma}'',\boldsymbol{\mu}'')$ *is rationally-compatible with* $(\boldsymbol{\lambda},\boldsymbol{\delta},\boldsymbol{\sigma},\boldsymbol{\mu})$.

Proof. We apply Corollary B.6 and obtain $(\delta'',\sigma'',\mu'')$. Let $\mathsf{inflow}''(Z)$ denote the quantity $\mathsf{inflow}(Z)$ measured on the vector $(\boldsymbol{\delta}'',\boldsymbol{\sigma}'',\boldsymbol{\mu}'')$.

$$\sum_{Z\subseteq V}\mathsf{inflow}''(Z)\cdot|Z|^2 \le \sum_{B\neq\emptyset}\lambda_B\cdot|B|^2 \tag{90}$$

Define the following quantities.

$$\begin{aligned}T_{\sigma''} &\stackrel{\text{def}}{=} \sum_{I\perp J}\sigma''_{I,J}(|I\cup J|^2+|I\cap J|^2-|I|^2-|J|^2)\\ T_{\delta''} &\stackrel{\text{def}}{=} \sum_{X\subset Y}\delta''_{Y|X}(|Y|^2-|X|^2)\\ T_{\mu''} &\stackrel{\text{def}}{=} \sum_{X\subset Y}\mu''_{X,Y}(|X|^2-|Y|^2)\end{aligned}$$

Now (90) becomes

$$T_{\sigma''} + T_{\delta''} + T_{\mu''} \le \sum_{B\neq\emptyset} \lambda_B \cdot |B|^2$$

$$T_{\sigma''} + T_{\delta''} - \sum_{Y\neq\emptyset} \mu''_{\emptyset,Y} \cdot |Y|^2 + \sum_{\emptyset\neq X\subset Y} \mu''_{X,Y}(|X|^2 - |Y|^2) \le \sum_{B\neq\emptyset} \lambda_B \cdot |B|^2$$

$$\begin{aligned} T_{\sigma''} + T_{\delta''} &\le \sum_{Y\neq\emptyset} \mu''_{\emptyset,Y} \cdot |Y|^2 + \sum_{B\neq\emptyset} |B|^2 \Bigg(\underbrace{\lambda_B - \sum_{Y\supset B} \mu''_{B,Y}}_{\ge 0 \text{ by } (86)} + \sum_{X\subset B} \mu''_{X,B} \Bigg) \\ &\le n^2 \cdot \underbrace{\left(\sum_{Y\neq\emptyset} \mu''_{\emptyset,Y} \right)}_{\text{bounded by } (87)} + n^2 \cdot \underbrace{\sum_{B\neq\emptyset} \left(\lambda_B - \sum_{Y\supset B} \mu''_{B,Y} + \sum_{X\subset B} \mu''_{X,B} \right)}_{=\|\boldsymbol{\lambda}\|_1} \\ &\le n^3 \cdot \|\boldsymbol{\lambda}\|_1 \end{aligned}$$

Claim 1: For any $I \perp J$, we have $|I \cup J|^2 + |I \cap J|^2 - |I|^2 - |J|^2 \ge 2$.

From Claim 1, $T_{\sigma''} \ge 2\|\boldsymbol{\sigma}\|_1$. Moreover, for any $X \subset Y$, $|Y|^2 - |X|^2 \ge 1$. Therefore, $T_{\delta''} \ge \|\boldsymbol{\delta}''\|_1$, which proves (88).

To prove Claim 1, let $a \stackrel{\text{def}}{=} |I \setminus J|$, $b \stackrel{\text{def}}{=} |J \setminus I|$, and $c = |I \cap J|$. Because $I \perp J$, we have $a \ge 1$ and $b \ge 1$.

$$|I \cup J|^2 + |I \cap J|^2 - |I|^2 - |J|^2 = (a+b+c)^2 + c^2 - (a+c)^2 - (b+c)^2 = 2ab \ge 2.$$

□

B.3 Construction of a poly-sized proof sequence

Given a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$, Section 5 shows that the runtime of PANDA algorithm depends on the length of the proof sequence that is being used for that inequality, thus motivating then need to minimize that length as much as possible. Section 4.1 shows that the inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ basically corresponds to a feasible solution to the linear program (42), whose size is $O(2^{2n})$. Our aim in this section is to construct a proof sequence for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ whose length is polynomial in the size of that linear program. Hence, we are looking for a proof sequence whose length is poly(2^n).

We start by introducing a connection to flow networks in Section B.3.1, and then we present the actual construction of a poly-sized proof sequence in Section B.3.2. The construction relies on many technical tools among which is Edmond-Karp maximum flow algorithm [18]. It also relies on the bounds developed earlier in Section B.2.

B.3.1 Introduction: Connection to flow networks

We will make use of the following definition. Given a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$, define

$$\mathsf{filter}(\boldsymbol{\lambda}) \stackrel{\text{def}}{=} \left\{ F \mid \exists B \in 2^V \text{ where } \lambda_B > 0 \text{ and } B \subseteq F \right\}. \tag{91}$$

Theorem B.8 (Construction of a proof sequence using a flow network)**.** *Let $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ be a Shannon flow inequality with witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$. Algorithm 2 produces a proof sequence for the inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with length at most $2^n D(\|\boldsymbol{\lambda}\|_1 + \|\boldsymbol{\sigma}\|_1)$, where D is the minimum common denominator of all entries in $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma})$.*

Algorithm 2 Constructing a proof sequence for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}, \mathbf{h} \rangle$

Input: Inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$

1: $\mathsf{ProofSeq} \leftarrow ()$
2: **For** each $B \in 2^V$ **do**
3: $\quad t \leftarrow \min\{\lambda_B, \delta_{B|\emptyset}\}; \quad \lambda_B \leftarrow \lambda_B - t; \quad \delta_{B|\emptyset} \leftarrow \delta_{B|\emptyset} - t$
4: $w \leftarrow 1/D$
5: **While** $\mathsf{filter}(\boldsymbol{\lambda}) \neq \emptyset$ **do**
6: $\quad \mathcal{K} \leftarrow \{Z \mid Z \text{ is reachable from } \emptyset \text{ in } G(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})\}$
7: $\quad$ **If** $(\mathsf{filter}(\boldsymbol{\lambda}) \cap \mathcal{K} \neq \emptyset)$ **then** ▷ Case 1
8: $\quad\quad$ Find a shortest path $\emptyset = X_0, X_1, \ldots, X_\ell = B$ in $G(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ from $\emptyset$ to some B where $\lambda_B > 0$
9: $\quad\quad$ **For** $(j \leftarrow 1$ **to** $\ell)$ **do** ▷ Push w-flow to $\mathbb{F}$
10: $\quad\quad\quad$ **If** $((X_{j-1}, X_j)$ is an up arc$)$ **then**
11: $\quad\quad\quad\quad$ Append $w \cdot \mathbf{c}_{X_{j-1}, X_j}$ to $\mathsf{ProofSeq}$; $\quad \delta_{X_j|X_{j-1}} \leftarrow \delta_{X_j|X_{j-1}} - w$
12: $\quad\quad\quad$ **else** ▷ (X_{j-1}, X_j) is a down arc
13: $\quad\quad\quad\quad$ Append $w \cdot \mathbf{d}_{X_{j-1}, X_j}$ to $\mathsf{ProofSeq}$; $\quad \delta_{X_{j-1}|X_j} \leftarrow \delta_{X_{j-1}|X_j} + w$
14: $\quad\quad \lambda_B \leftarrow \lambda_B - w$;
15: $\quad$ **else If** (There exists a good pair (I, J) for $\mathcal{K}$) **then** ▷ Case 2
16: $\quad\quad$ Find a shortest path $\emptyset = X_0, X_1, \ldots, X_\ell = I$
17: $\quad\quad$ **For** $(j \leftarrow 1$ **to** $\ell)$ **do** ▷ Push w-flow to I
18: $\quad\quad\quad$ **If** $((X_{j-1}, X_j)$ is an up arc$)$ **then**
19: $\quad\quad\quad\quad$ Append $w \cdot \mathbf{c}_{X_{j-1}, X_j}$ to $\mathsf{ProofSeq}$; $\quad \delta_{X_j|X_{j-1}} \leftarrow \delta_{X_j|X_{j-1}} - w$
20: $\quad\quad\quad$ **else** ▷ (X_{j-1}, X_j) is a down arc
21: $\quad\quad\quad\quad$ Append $w \cdot \mathbf{d}_{X_{j-1}, X_j}$ to $\mathsf{ProofSeq}$; $\quad \delta_{X_{j-1}|X_j} \leftarrow \delta_{X_{j-1}|X_j} + w$
22: $\quad\quad$ Append $w \cdot \mathbf{d}_{I, I \cap J}$, then $w \cdot \mathbf{s}_{I,J}$ to $\mathsf{ProofSeq}$
23: $\quad\quad \sigma_{I,J} \leftarrow \sigma_{I,J} - w; \quad \delta_{I \cup J|J} \leftarrow \delta_{I \cup J|J} + w; \quad \delta_{I \cap J|\emptyset} \leftarrow \delta_{I \cap J|\emptyset} + w.$
24: **Return** $\mathsf{ProofSeq}$

Proof. Given $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$, define a flow network $G(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu}) = (2^V, \mathcal{A})$ as follows. There is an arc (i.e. directed edge) $(X, Y) \in \mathcal{A}$ for every pair (X, Y) for which $X \subset Y$ and $\delta_{Y|X} > 0$; the capacity of (X, Y) is $\delta_{Y|X}$. These are called *up arcs*. There is an arc (Y, X) for every pair (X, Y) such that $X \subset Y$; The capacity of this arc is $+\infty$. These are called *down arcs*. Let $\mathcal{K}$ denote the set of vertices Z reachable from $\emptyset$ in $G(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$. A pair $(I, J) \in 2^V \times 2^V$ is *good* for $\mathcal{K}$ if $I \subseteq \mathcal{K}$, $J \subseteq \mathcal{K}$, $I \cup J \not\subseteq \mathcal{K}$, and $\sigma_{I,J} > 0$.

For any set $\mathcal{K} \subset 2^V$ such that $\mathsf{filter}(\boldsymbol{\lambda}) \cap \mathcal{K} = \emptyset$, define $\mathsf{outflow}(\mathcal{K}) \stackrel{\text{def}}{=} \sum_{Z \notin \mathcal{K}} \mathsf{inflow}(Z)$. We claim that the algorithm maintains the following two invariants:

- **Invariant 1:** all the quantities $\mathsf{inflow}(Z) - \lambda_Z$ are unchanged from one iteration to the next iteration of the algorithm.
- **Invariant 2:** for every $\mathcal{K}$ where $\mathcal{K} \cap \mathsf{filter}(\boldsymbol{\lambda}) = \emptyset$, at the beginning of each iteration we have

$$\sum_{B \in \mathsf{filter}(\boldsymbol{\lambda})} \lambda_B \leq \mathsf{outflow}(\mathcal{K}).$$

The first invariant is simple to verify: we specifically modified all the entries $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ to keep the flows unchanged, except for the very last element B in line 8, where the value $\mathsf{inflow}(B)$ for that element B is reduced by w. However, λ_B is also reduced by w at line 14, keeping $\mathsf{inflow}(B) - \lambda_B$ constant. The second invariant is satisfied at the beginning of the very first iteration, and hence it is satisfied the the beginning of every later iteration due to the first invariant.

We now show that the algorithm produces a valid proof sequence. At the beginning of every iteration, the inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ is a Shannon flow inequality witnessed by $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, due to invariant 1 and

Proposition 4.3. (Note that we modify $\boldsymbol{\lambda}$ and the witness inside each iteration.) Hence, if we can show that the algorithm does terminate, then the proof sequence it produces is valid. To show termination, we show that as long as there exists a $B \in \mathsf{filter}(\boldsymbol{\lambda})$ for which $\lambda_B > 0$, then either Case 1 or Case 2 applies in the algorithm.

Assume that at the beginning of some iteration, we have $\mathsf{filter}(\boldsymbol{\lambda}) \cap \mathcal{K} \neq \emptyset$. Then from (91), there is some $B \subseteq F$ where $\lambda_B > 0$ and $F \in \mathcal{K}$, and thanks to the down arc (F, B), we have $B \in \mathcal{K}$. Now, assume that at the beginning of some iteration (where $\lambda_B > 0$ for some B), we have $\mathsf{filter}(\boldsymbol{\lambda}) \cap \mathcal{K} = \emptyset$, yet there is no good pair (I, J) for $\mathcal{K}$. Then, $\mathsf{outflow}(\mathcal{K}) \geq \lambda_B > 0$ due to Invariant 2. However, all variables $\delta_{Y|X}, \mu_{X,Y}, \sigma_{I,J}$ contribute a non-positive amount to $\mathsf{outflow}(\mathcal{K})$, which is a contradiction.

Next, we bound the proof sequence's length. When any one of the two cases applies, $\|\boldsymbol{\lambda}\|_1 + \|\boldsymbol{\sigma}\|_1$ was reduced by $w = 1/D$, and at most 2^n steps are added to the proof sequence. Hence the overall length of the proof sequence is at most $2^n D(\|\boldsymbol{\lambda}\|_1 + \|\boldsymbol{\sigma}\|_1)$. □

B.3.2 Final construction of a poly-sized proof sequence

We start by extending the flow network that we constructed earlier in Theorem B.8. Then, we prove a lower bound on the maximum flow in Lemma B.10. Theorem B.12 presents an advanced construction of a proof sequence that is "almost" polynomial. Among other technical tools, the construction uses Edmond-Karp maximum flow Algorithm as a black-box [18]. Finally, Corollary B.14 makes the final step and bounds the length of the proof sequence to be truly polynomial in $O(2^n)$.

Definition B.9 (Extended flow network)**.** Given $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$, we define a flow network $\bar{G}(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ (which is an extended version of the network $G(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ defined in the proof of Theorem B.8) as follows. The set of vertices of the network $\bar{G}$ is $\bar{\mathcal{V}} \stackrel{\text{def}}{=} 2^V \cup \mathcal{T} \cup \{\bar{T}\}$, where $\mathcal{T}$ is a set that contains a new vertex $T_{I,J}$ for every pair $I, J \subseteq V$ such that $I \perp J$, i.e.

$$\mathcal{T} \stackrel{\text{def}}{=} \{T_{I,J} \mid I, J \subseteq V \wedge I \perp J\},$$

and $\bar{T}$ is yet another vertex that represents the *sink* of the network. The *source* of the network $\bar{G}$ is $\emptyset$. The set of arcs $\bar{\mathcal{A}}$ of $\bar{G}$ consists of the following subsets:

- There is an arc (X, Y) for every pair (X, Y) for which $X \subset Y \subseteq V$ and $\delta_{Y|X} > 0$; the capacity of (X, Y) is $\delta_{Y|X}$. These are called *up arcs* (as in the network G from Theorem B.8).
- There is an arc (Y, X) for every pair (X, Y) such that $X \subset Y \subseteq V$; The capacity of this arc is $+\infty$. These are called *down arcs* (as in the network G).
- For ever pair $I, J \subseteq V$ where $I \perp J, \sigma_{I,J} > 0$, there are two arcs $(I, T_{I,J})$ and $(J, T_{I,J})$ each of which has an infinite capacity $(+\infty)$, and there is a third arc $(T_{I,J}, \bar{T})$ whose capacity is $\sigma_{I,J}$.
- For every $\emptyset \neq B \subseteq V$ where $\lambda_B > 0$, there is an arc $(B, \bar{T})$ whose capacity is λ_B.

Lemma B.10 (Maximum flow is $\geq \|\boldsymbol{\lambda}\|_1$)**.** *For any $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$, the maximum flow of the network $\bar{G}(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$ given by Definition B.9 is $\geq \|\boldsymbol{\lambda}\|_1$.*

Proof. We will use the min-cut max-flow theorem, and show that every cut has a capacity $\geq \|\boldsymbol{\lambda}\|_1$ (where the capacity of a cut is the total capacity of arcs crossing that cut). In particular, for any $\mathcal{C} \subset \bar{\mathcal{V}}$ such that $\emptyset \in \mathcal{C}, \bar{T} \notin \mathcal{C}$, and $\mathcal{C}' \stackrel{\text{def}}{=} \bar{\mathcal{V}} \setminus \mathcal{C}$, we will show that the capacity of the $\mathcal{C}$-$\mathcal{C}'$ cut is $\geq \|\boldsymbol{\lambda}\|_1$.

- If there are $X \subset Y \subseteq V$ where $Y \in \mathcal{C}$ and $X \notin \mathcal{C}$, then the down arc (Y, X) crosses the cut, hence the cut capacity is $+\infty$.
- Otherwise (i.e. if for any $X \subset Y \subseteq V$ where $Y \in \mathcal{C}$, we have $X \in \mathcal{C}$), if there is a pair $I, J \subseteq V$ where $I \perp J$, $\sigma_{I,J} > 0$, $I \in \mathcal{C}$ and $T_{I,J} \notin \mathcal{C}$, then the cut capacity is also $+\infty$.

- Otherwise, let $\mathcal{K} \stackrel{\text{def}}{=} \mathcal{C} \cap 2^V$, $\mathcal{K}' \stackrel{\text{def}}{=} 2^V \setminus \mathcal{K}$.

$$\begin{aligned} \sum_{Z \in \mathcal{K}'} \mathsf{inflow}(Z) &\geq \sum_{B \in \mathcal{K}'} \lambda_B \\ \sum_{\substack{I,J \in \mathcal{K} \\ I \perp J,\ I \cup J \in \mathcal{K}'}} \sigma_{I,J} + \sum_{\substack{X \in \mathcal{K},\ Y \in \mathcal{K}' \\ X \subset Y}} \delta_{Y|X} &\geq \sum_{B \in \mathcal{K}'} \lambda_B \end{aligned}$$

  However, we have 3 types of arcs crossing the cut: $(T_{I,J}, \bar{T})$ where $I \in \mathcal{K}$, up arcs (X, Y) where $X \in \mathcal{K}, Y \in \mathcal{K}'$, and $(B, \bar{T})$ where $B \in \mathcal{K}$. Hence the cut capacity is at least

$$\sum_{\substack{I \in \mathcal{K} \\ I \perp J}} \sigma_{I,J} + \sum_{\substack{X \in \mathcal{K},\ Y \in \mathcal{K}' \\ X \subset Y}} \delta_{Y|X} + \sum_{B \in \mathcal{K}} \lambda_B \geq \|\boldsymbol{\lambda}\|_1.$$

□

Lemma B.11 (Getting rid of the domination assumption)**.** *Let $\langle \boldsymbol{\delta}'_{\ell'}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}'_0, \mathbf{h} \rangle$ be an equality that has a proof sequence $\mathsf{ProofSeq}'$ of length ℓ', and let $\boldsymbol{\delta}_0$ be a vector that dominates $\boldsymbol{\delta}'_0$. Then, there exists a vector $\boldsymbol{\delta}_\ell$ that dominates $\boldsymbol{\delta}'_{\ell'}$ and an inequality $\langle \boldsymbol{\delta}_\ell, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}_0, \mathbf{h} \rangle$ that has a proof sequence $\mathsf{ProofSeq}$ of length ℓ that satisfies*

$$\ell = O(3^n \cdot \ell').$$

Proof. Let f be the domination function of $(\boldsymbol{\delta}_0, \boldsymbol{\delta}'_0)$. Let $w' \cdot \mathbf{f}' = \boldsymbol{\delta}'_1 - \boldsymbol{\delta}'_0$ be the first proof step in $\mathsf{ProofSeq}'$, where $w' > 0$ and $\mathbf{f}' \in \{\mathbf{s}_{I,J}, \mathbf{m}_{X,Y}, \mathbf{c}_{X,Y}, \mathbf{d}_{Y,X}\}$. Initially, let $\boldsymbol{\delta}_1$ be identical to $\boldsymbol{\delta}_0$.

- If $\mathbf{f}' = \mathbf{s}_{I,J}$ for some $I \perp J$, then $(I \cap J, I)$ is dominated by a total of w' (i.e. $\sum_{(X,Y)} f((X,Y),(I \cap J, I)) \geq w'$). We keep looking for pairs $(X, Y) \in \mathcal{P}$ that dominate $(I \cap J, I)$ (i.e. such that $f((X,Y),(I \cap J, I)) > 0$), and we make those pairs (X, Y) dominate $(J, I \cup J)$ instead (i.e. we reduce $f((X,Y),(I \cap J, I))$ and increase $f((X,Y),(J, I \cup J))$ by the same amount). We keep doing so until $(J, I \cup J)$ is dominated by a total of w'. Now, f is a domination function for $(\boldsymbol{\delta}_1, \boldsymbol{\delta}'_1)$.
- If $\mathbf{f}' = \mathbf{m}_{X',Y'}$ for some $X' \subset Y'$, then we look for $(\emptyset, Y)$ that dominate $(\emptyset, Y')$ and we make them dominate $(\emptyset, X')$ instead until $(\emptyset, X')$ is dominated by a total of w'. Now, f is a domination function for $(\boldsymbol{\delta}_1, \boldsymbol{\delta}'_1)$.
- If $\mathbf{f}' = \mathbf{d}_{Y',X'}$ for some $X' \subset Y'$, then we look for $(\emptyset, Y)$ that dominate $(\emptyset, Y')$, and we append $w \cdot \mathbf{m}_{Y',Y}$ to $\mathsf{ProofSeq}$ (if $Y' \neq Y$), and reduce $f((\emptyset, Y),(\emptyset, Y'))$ by the same amount w. We keep doing so until the total of w is equal to w'. Now, we append $w' \cdot \mathbf{d}_{X',Y'}$ to $\mathsf{ProofSeq}$, and increase $f((\emptyset, X'),(\emptyset, X'))$ and $f((X', Y'),(X', Y'))$ by w'.
- If $\mathbf{f}' = \mathbf{c}_{X',Y'}$ for some $X' \subset Y'$, then we look for $(\emptyset, X)$ that dominate $(\emptyset, X')$, and we append $w \cdot \mathbf{m}_{X',X}$ to $\mathsf{ProofSeq}$, and reduce $f((\emptyset, X),(\emptyset, X'))$ by the same amount w, until the total of w is equal to w'. Now, we look for (X, Y) that dominates (X', Y') [16], and append $w_2 \cdot \mathbf{m}_{A,X'}$, (where $A \stackrel{\text{def}}{=} X' \setminus (Y \setminus X)$) $w_2 \cdot \mathbf{s}_{Y,A}$, $w_2 \cdot \mathbf{c}_{A,Y \cup A}$, $w_2 \cdot \mathbf{m}_{Y',Y \cup A}$ to $\mathsf{ProofSeq}$, and we reduce $f((X,Y),(X',Y'))$ by the same amount w_2. We keep doing so until the total of w_2 is equal to w'. Finally, we increase $f((\emptyset, Y'),(\emptyset, Y'))$ by w'. The total number of pairs (X, Y) where $X \subset Y \subseteq V$ is $\leq 3^n$.

□

[16] From Definition B.1, if $f((X,Y),(X',Y')) > 0$, then $Y' \setminus X' \subseteq Y \setminus X$ and $X' \supseteq X$.

Algorithm 3 Constructing a poly-sized proof sequence for $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$

Input: Inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$, where $\|\boldsymbol{\lambda}\|_1 = 1$

$\mathsf{ProofSeq} \leftarrow ()$
$D \leftarrow$ LEAST-COMMON-DENOMINATOR$((\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu}))$ ▷ D is the least common denominator of all entries
$\boldsymbol{\lambda} \leftarrow D \cdot \boldsymbol{\lambda}$; $\boldsymbol{\delta} \leftarrow D \cdot \boldsymbol{\delta}$; $\boldsymbol{\sigma} \leftarrow D \cdot \boldsymbol{\sigma}$; $\boldsymbol{\mu} \leftarrow D \cdot \boldsymbol{\mu}$ ▷ Now, all entries are integers
Find $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ with witness $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ where $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$ and $\|\boldsymbol{\sigma}'\|_1 \le \frac{1}{2} n^3 \|\boldsymbol{\lambda}\|_1$
▷ Corollary B.7
$\Delta \leftarrow \max\{2^i \mid 2^i \le \|\boldsymbol{\lambda}\|_1 \text{ and } i \in \mathbb{N}\}$ ▷ Now, we have $\Delta \le \|\boldsymbol{\lambda}\|_1 < 2\Delta$ and $\|\boldsymbol{\sigma}'\|_1 < n^3 \Delta$
While $\|\boldsymbol{\lambda}\|_1 > 0$ **do**
  Construct $\bar{G}(\boldsymbol{\lambda}, \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$ ▷ Definition B.9
  $P \leftarrow$EDMOND-KARP$(\bar{G}(\boldsymbol{\lambda}, \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}'))$. ▷ Apply Edmond-Karp algorithm for maximum flow
  ▷ P is a collection of augmenting paths
  For each augmenting path $\emptyset = X_0, X_1, \ldots, X_\ell = \bar{T}$ with capacity c in P **do**
    For $j \leftarrow 1$ to ℓ **do**
      If $(X_{j-1}, X_j) = (B, \bar{T})$ for some $\emptyset \ne B \subseteq V, \lambda_B > 0$ **then**
        $\lambda_B \leftarrow \lambda_B - c$; $\delta'_{B|\emptyset} \leftarrow \delta'_{B|\emptyset} - c$
        $b \leftarrow c$
        While $b > 0$ **do**
          Find $Y \supseteq B$ such that $\delta_{Y|\emptyset} > 0$ ▷ Possible because $\boldsymbol{\delta}$ dominates $\boldsymbol{\delta}'$
          $t \leftarrow \min\{b, \delta_{Y|\emptyset}\}$; $b \leftarrow b - t$; $\delta_{Y|\emptyset} \leftarrow \delta_{Y|\emptyset} - t$
          If $Y \supset B$ **then**
            Append $b \cdot \mathbf{d}_{Y,B}$ to $\mathsf{ProofSeq}$; $\delta_{Y|B} \leftarrow \delta_{Y|B} + t$
      else
        $\boldsymbol{\delta}'_0 \leftarrow \boldsymbol{\delta}'$; $\mathsf{ProofSeq}'_0 \leftarrow ()$; $\boldsymbol{\delta}_0 \leftarrow \boldsymbol{\delta}$
        If (X_{j-1}, X_j) is an up arc **then**
          Append $c \cdot \mathbf{c}_{X_{j-1}, X_j}$ to $\mathsf{ProofSeq}'_0$; $\boldsymbol{\delta}' \leftarrow \boldsymbol{\delta}' + c \cdot \mathbf{c}_{X_{j-1}, X_j}$
        else If (X_{j-1}, X_j) is a down arc **then**
          Append $c \cdot \mathbf{d}_{X_{j-1}, X_j}$ to $\mathsf{ProofSeq}'_0$; $\boldsymbol{\delta}' \leftarrow \boldsymbol{\delta}' + c \cdot \mathbf{d}_{X_{j-1}, X_j}$
        else If $(X_{j-1}, X_j) = (I, T_{I,J})$ for some $I, J \subseteq V, I \perp J, \sigma_{I,J} > 0$ **then**
          Append $c \cdot \mathbf{d}_{I, I \cap J}$, then $c \cdot \mathbf{s}_{I,J}$ to $\mathsf{ProofSeq}'_0$; $\boldsymbol{\delta}' \leftarrow \boldsymbol{\delta}' + c \cdot \mathbf{d}_{I, I\cap J} + c \cdot \mathbf{s}_{I,J}$
          $\sigma'_{I,J} \leftarrow \sigma'_{I,J} - c$
            ▷ Now, $\mathsf{ProofSeq}'_0$ is a proof sequence for $\langle \boldsymbol{\delta}', \mathbf{h} \rangle \le \langle \boldsymbol{\delta}'_0, \mathbf{h} \rangle$, and $\boldsymbol{\delta}_0$ dominates $\boldsymbol{\delta}'_0$
        Find $\boldsymbol{\delta}$ dominating $\boldsymbol{\delta}'$ and a proof sequence $\mathsf{ProofSeq}_0$ for $\langle \boldsymbol{\delta}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}_0, \mathbf{h} \rangle$ ▷ Lemma B.11
        Append $\mathsf{ProofSeq}_0$ to $\mathsf{ProofSeq}$
  If $\|\boldsymbol{\lambda}\|_1 < \Delta$ **then**
    Find $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}', \mathbf{h} \rangle$ with witness $(\boldsymbol{\sigma}', \boldsymbol{\mu}')$ where
      $\boldsymbol{\delta}'$ is dominated by $\boldsymbol{\delta}$ and $\|\boldsymbol{\sigma}'\|_1 \le \frac{1}{2} n^3 \|\boldsymbol{\lambda}\|_1$ ▷ Corollary B.7
    $\Delta \leftarrow \max\{2^i \mid 2^i \le \|\boldsymbol{\lambda}\|_1 \text{ and } i \in \mathbb{N}\}$ ▷ Now, we have $\Delta \le \|\boldsymbol{\lambda}\|_1 < 2\Delta$ and $\|\boldsymbol{\sigma}'\|_1 < n^3 \Delta$
Return $\mathsf{ProofSeq}$

Theorem B.12 (Construction of a poly-sized proof sequence, modulo $\log D$)**.** *For any Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \le \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ with witness $(\boldsymbol{\sigma}, \boldsymbol{\mu})$ where $\|\boldsymbol{\lambda}\|_1 = 1$, Algorithm 3 produces a proof sequence whose length is $O(\log D \cdot n^3 \cdot 2^{6n} \cdot 3^n) = O(\log D \cdot \mathrm{poly}(2^n))$, where D is the common denominator of all entries in $(\boldsymbol{\lambda}, \boldsymbol{\delta}, \boldsymbol{\sigma}, \boldsymbol{\mu})$.*

Proof. Given a flow network G with vertex set $\mathcal{V}$ and arc set $\mathcal{A}$, Edmond-Karp algorithm [15, 18] finds the maximum flow in time $O(|\mathcal{V}| \cdot |\mathcal{A}|^2)$. Edmond-Karp is based on augmenting paths. In particular, it finds $O(|\mathcal{V}| \cdot |\mathcal{A}|)$ such paths, each of which has length $\le |\mathcal{V}|$ and can be found in time $O(|\mathcal{A}|)$. Each path also has a *capacity* c, which is the minimum capacity among arcs of that path. The total capacity of all augmenting paths is equal to the maximum flow.

The flow network $\bar{G}(\boldsymbol{\lambda}, \boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$ constructed in Line 8 of Algorithm 3 has a vertex set $\mathcal{V}$ of size $|\bar{\mathcal{V}}| = O(2^{2n})$ and an arc set $\mathcal{A}$ of size $|\bar{\mathcal{A}}| = O(2^{2n})$. By Lemma B.10, it has a maximum flow $\geq \|\boldsymbol{\lambda}\|_1 \geq \Delta$. The loop in Line 11 maintains the quantities $\mathsf{inflow}'(Z) - \lambda_Z$ for all $\emptyset \neq Z \subseteq V$ unchanged (where $\mathsf{inflow}'(Z)$ is the quantity $\mathsf{inflow}(Z)$ measured on the vector $(\boldsymbol{\delta}', \boldsymbol{\sigma}', \boldsymbol{\mu}')$.). For each augmenting path in the loop, we either reduce some λ_B (in Line 14) or $\sigma'_{I,J}$ (in Line 29). After we are done with all paths in the loop, the quantity $\|\boldsymbol{\sigma}'\|_1 + \|\boldsymbol{\lambda}\|_1$ is reduced by at least Δ, and a total of $O(2^{6n})$ proof steps where appended to $\mathsf{ProofSeq}'_0$. By Lemma B.11, the original proof sequence $\mathsf{ProofSeq}$ is extended by $O(2^{6n} \cdot 3^n)$ proof steps.

Before the loop in line 7, we had $\Delta \leq \|\boldsymbol{\lambda}\|_1 < 2\Delta$ and $\|\boldsymbol{\sigma}'\|_1 < n^3\Delta$, which both continue to hold and Δ remains fixed as long as $\|\boldsymbol{\lambda}\|_1$ remains $\geq \Delta$ in line 32. The maximum number of iterations the loop in line 7 can make before $\|\boldsymbol{\lambda}\|_1$ drops below Δ can be bounded as follows. At each iteration, the quantity $\|\boldsymbol{\sigma}'\|_1 + \|\boldsymbol{\lambda}\|_1$ is reduced by at least Δ. However, $\|\boldsymbol{\lambda}\|_1$ cannot be reduced by more than Δ in total (since $\Delta \leq \|\boldsymbol{\lambda}\|_1 < 2\Delta$). Also, $\|\boldsymbol{\sigma}'\|_1$ cannot be reduced by more than $n^3\Delta$ in total. Therefore the maximum number of iterations before $\|\boldsymbol{\lambda}\|_1$ drops below Δ is bounded by $n^3 + 1$.

Every time $\|\boldsymbol{\lambda}\|_1$ drops below Δ, Δ gets divided by 2 (or a higher power of 2). The initial Δ in line 6 was $\leq D$. Therefore, the maximum number of times $\|\boldsymbol{\lambda}\|_1$ can drop below Δ is $O(\log D)$. □

Proposition B.13 ($\log D$ is polynomial)**.** *Given an optimization problem of the form (38), if the objective value is positive and bounded, then there exists a vector $\boldsymbol{\lambda} = (\lambda_B)_{B \in \mathcal{B}}$ satisfying the following conditions:*

(a) $\|\boldsymbol{\lambda}\|_1 = 1$.

(b) The optimization problem (38) has the same optimal objective value as the linear program (41) (using this $\boldsymbol{\lambda}$).

(c) The linear program (41) has an optimal dual solution $(\boldsymbol{\delta}^*, \boldsymbol{\sigma}^*, \boldsymbol{\mu}^*)$ such that the common denominator D of all entries in $(\boldsymbol{\lambda}, \boldsymbol{\delta}^*, \boldsymbol{\sigma}^*, \boldsymbol{\mu}^*)$ satisfies

$$D \leq (2^n)!$$

Proof. This proposition is a minor extension to Lemma 4.2, specialized to the optimization problem (38). In particular, when we pick an optimal dual solution $(\mathbf{z}^*, \mathbf{y}^*)$ to (82), we choose $(\mathbf{z}^*, \mathbf{y}^*)$ to be an extreme point of the following polyhedron

$$\{(\mathbf{z}, \mathbf{y}) \mid \mathbf{A}^T\mathbf{y} \geq \mathbf{C}\mathbf{z}, \mathbf{1}_p^T\mathbf{z} \geq 1, (\mathbf{z}, \mathbf{y}) \geq \mathbf{0}\}.$$

In (38), all entries of $\mathbf{A}$ and $\mathbf{C}$ are in $\{1, 0, -1\}$. By Cramer's rule, the common denominator D^* of all entries in $(\mathbf{z}^*, \mathbf{y}^*)$ is (the absolute value of) the determinant of an $m \times m$ matrix whose entries are all in $\{1, 0, -1\}$. Hence, $D \leq m!$. □

Corollary B.14 (Construction of a poly-sized proof sequence)**.** *Given any optimization problem of the form (38) where the optimal objective value OBJ is positive and bounded, there exists a Shannon flow inequality $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ satisfying the following conditions:*

- $\sum_{(X,Y)} \delta_{Y|X} n_{Y|X} \leq \mathsf{OBJ}$.
- $\lambda_{Y|X} = 0$ *for all* $(X, Y) \in \mathcal{P}$ *where* $X \neq \emptyset$ *or* $Y \notin \mathcal{B}$.
- $\langle \boldsymbol{\lambda}, \mathbf{h} \rangle \leq \langle \boldsymbol{\delta}, \mathbf{h} \rangle$ *has a proof sequence of length* $O(n^4 \cdot 2^{7n} \cdot 3^n) = O(\mathrm{poly}(2^n))$.

C More details on Section 5

Proof of Lemma 5.3. To obtain the sub-tables $R^{(j)}$, observe that the number of tuples $\mathbf{t} \in \Pi_X(R)$ with log-degree in the interval $[j, j+1)$ is at most $|R|/2^j \leq 2^{n_{Y|\emptyset} - j}$. Hence, if we partition R based on which of the buckets $[j, j+1)$ the log-degree falls into, we would almost have the required inequality: $n^{(j)}_{Y|X} + n^{(j)}_{X|\emptyset} \leq (j+1) + (n_{Y|\emptyset} - j) = n_{Y|\emptyset} + 1$. To resolve the situation, we partition each $R^{(j)}$ into two equal-sized tables. Overall, we need $k = 2 \log |R|$. □

D More details on Section 6

Proof of Propositions 3.4 and 6.3. To prove (31) and (56), note that for $h \in \mathsf{SA}_n$ and any set $F \subseteq V$, from subadditivity we have $h(F) \leq \sum_{v \in F} h(v) \leq |F|$. Recall the definition of tree-width in (16), we have

$$\mathsf{Minimaxwidth}_{\mathsf{SA}_n \cap \mathsf{VD}(\mathbf{1}_V)}(Q) = \min_{(T,\chi)} \max_{t \in V(T)} \sup_{h \in \mathsf{SA}_n \cap \mathsf{VD}(\mathbf{1}_V)} h(\chi(t)) \leq \min_{(T,\chi)} \max_{t \in V(T)} |\chi(t)| = \mathsf{tw}(\mathcal{H}) + 1.$$

To show the reverse, define the function $\bar{h}(F) = |F|$ for all $F \subseteq V$. This function is modular and vertex-dominated, and thus

$$\begin{aligned}
\mathsf{tw}(\mathcal{H}) + 1 &\geq \mathsf{Minimaxwidth}_{\mathsf{SA}_n \cap \mathsf{VD}(\mathbf{1}_V)}(Q) \\
&\geq \mathsf{Minimaxwidth}_{\mathsf{M}_n \cap \mathsf{VD}(\mathbf{1}_V)}(Q) \\
&\geq \mathsf{Maximinwidth}_{\mathsf{M}_n \cap \mathsf{VD}(\mathbf{1}_V)}(Q) \\
&= \sup_{h \in \mathsf{M}_n \cap \mathsf{VD}(\mathbf{1}_V)} \min_{(T,\chi)} \max_{t \in V(T)} h(\chi(t)) \\
&\geq \min_{(T,\chi)} \max_{t \in V(T)} \bar{h}(\chi(t)) \\
&= \mathsf{tw}(\mathcal{H}) + 1.
\end{aligned}$$

Thus the bound hierarchy collapses for the $\mathsf{VD}(\mathbf{1}_V)$ constraints. Identity (31) is proved in the exact same way.

Identities (32) and the first equality in (57) are proved similarly, taking into account the proof of Proposition 2.7. To prove the second equality in (57), define the function $\bar{h}(B) = \rho(B)$ for all $B \subseteq V$. This function is in SA_n and is also edge-dominated. Then

$$\begin{aligned}
\mathsf{Maximinwidth}_{\mathsf{SA}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})}(Q) &= \sup_{h \in \mathsf{SA}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})} \min_{(T,\chi)} \max_{t \in \chi(t)} h(\chi(t)) \\
&\geq \min_{(T,\chi)} \max_{t \in \chi(t)} \bar{h}(\chi(t)) \\
&= \mathsf{ghtw}(\mathcal{H}) \\
&= \mathsf{Minimaxwidth}_{\mathsf{SA}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})}(Q).
\end{aligned}$$

We next prove (33), (34), and (58). Fix $B \subseteq V$ and let $h^* = \operatorname{argmax}\{h(B) \mid h \in \Gamma_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})\}$. Up to renumbering V, we can assume $B = [k]$ for some positive integer k, and $V = [n]$. By convention we have $[0] = \emptyset$. Define a set function $\bar{h}$ by setting $\bar{h}(F) = \sum_{i \in F}(h^*([i]) - h^*([i-1]))$. (And $\bar{h}(\emptyset) = 0$.) Clearly $\bar{h}$ is modular with $\bar{h}(B) = h^*(B)$. If we can show that $\bar{h}(F) \leq h^*(F)$ for all $F \subseteq V$, then $\bar{h} \in \mathsf{M}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})$, which implies

$$\max\{h(B) \mid h \in \mathsf{M}_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})\} = \max\{h(B) \mid h \in \Gamma_n \cap \mathsf{ED}(\mathbf{1}_\mathcal{E})\}. \tag{92}$$

We prove $\bar{h}(F) \leq h^*(F)$ by induction on $|F|$. The base case when $|F| = 0$ is trivial. For the inductive step, let j be the maximum integer in F, then by noting that $|F \cap [j-1]| < |F|$, we have

$$\begin{aligned}
\bar{h}(F) &= h^*([j]) - h^*([j-1]) + \sum_{i \in F - \{j\}} (h^*([i]) - h^*([i-1])) \\
&= h^*([j]) - h^*([j-1]) + \bar{h}(F \cap [j-1]) \\
&= h^*(F \cup [j-1]) - h^*([j-1]) + \bar{h}(F \cap [j-1]) \\
\text{(induction hypothesis)} \quad &\leq h^*(F \cup [j-1]) - h^*([j-1]) + h^*(F \cap [j-1]) \\
\text{(submodularity of } h^*) \quad &\leq h^*(F).
\end{aligned}$$

To see (34), we apply (92) with $B = V$. The left-hand side is precisely the dual of the fractional edge cover LP and thus it is equal to $\mathsf{AGM}(Q)$. Identity (33) follows from (34). Identity (58) follows by picking B to be the bag with the worst-case bound for any tree decomposition.

Identities (59) and (60) are just their definitions. □

Proof of Lemma 6.2. We prove the only non-trivial inequality that $\mathsf{Maximinwidth}_{\mathcal{F}}(\mathcal{H}) \leq \mathsf{Minimaxwidth}_{\mathcal{F}}(\mathcal{H})$. From definition:

$$\begin{aligned}
\mathsf{Maximinwidth}_{\mathcal{F}}(\mathcal{H}) &= \sup_{h\in\mathcal{F}} \min_{(T,\chi)} \max_{t\in V(T)} h(\chi(t)) \\
&\leq \min_{(T,\chi)} \sup_{h\in\mathcal{F}} \max_{t\in V(T)} h(\chi(t)) \\
&\leq \min_{(T,\chi)} \max_{t\in V(T)} \sup_{h\in\mathcal{F}} h(\chi(t)) \\
&= \mathsf{Minimaxwidth}_{\mathcal{F}}(\mathcal{H}).
\end{aligned}$$

□

Proof of Proposition 6.7. Let ZY denote the Zhang-Yeung query defined in the proof of Theorem 3.1. Add one additional relation R_V to the query, with very large relation size bound N_V. Call the resulting query ZY^+. Then, for this query $\mathsf{eda\text{-}subw}(\mathsf{ZY}^+) = \mathsf{eda\text{-}fhtw}(\mathsf{ZY}^+) < \mathsf{da\text{-}subw}(\mathsf{ZY}^+) = \mathsf{da\text{-}fhtw}(\mathsf{ZY}^+)$. In fact, with the gap amplification trick we showed in the proof of Theorem 3.1, we can show that there are queries for which the gap is as large as one wants.

On the other hand, for the 4-cycle query C_4, we already showed in Example 6.9 that $2\log N = \mathsf{eda\text{-}fhtw}(C_4) > 3/2\log N = \mathsf{da\text{-}subw}(C_4)$. □

Proof of Lemma 6.12. For any $a \in A$, define $\beta^*(a) \stackrel{\text{def}}{=} \text{argmax}_{b\in B} f(a,b)$.

$$\min_{a\in A}\max_{b\in B} f(a,b) = \min_{a\in A} f(a,\beta^*(a)) \leq \max_{\beta\in B^A}\min_{a\in A} f(a,\beta(a)).$$

Conversely,

$$\max_{\beta\in B^A}\min_{a\in A} f(a,\beta(a)) \leq \max_{\beta\in B^A}\min_{a\in A} f(a,\beta^*(a)) = \min_{a\in A} f(a,\beta^*(a)) = \min_{a\in A}\max_{b\in B} f(a,b).$$

□